Submitted to the Annals of Statistics
arXiv: math.PR/0000000

arXiv:1512.02174v2 [stat.ME] 8 Dec 2015

HIGHER ORDER ESTIMATING EQUATIONS FOR HIGH-DIMENSIONAL MODELS

By James Robins, Lingling Li, Rajarshi Mukherjee, Eric Tchetgen Tchetgen and Aad van der Vaart

Harvard University and Universiteit Leiden

We introduce a new method of estimation of parameters in semiparametric and nonparametric models. The method is based on estimating equations that are U-statistics in the observations. The U-statistics are based on higher order influence functions that extend ordinary linear influence functions of the parameter of interest, and represent higher derivatives of this parameter. For parameters for which the representation cannot be perfect the method leads to a bias-variance trade-off, and results in estimators that converge at a slower than $\sqrt{n}$-rate. In a number of examples the resulting rate can be shown to be optimal. We are particularly interested in estimating parameters in models with a nuisance parameter of high dimension or low regularity, where the parameter of interest cannot be estimated at $\sqrt{n}$-rate, but we also consider efficient $\sqrt{n}$-estimation using novel nonlinear estimators. The general approach is applied in detail to the example of estimating a mean response when the response is not always observed.

1. Introduction. Let $X_1, X_2, \ldots, X_n$ be a random sample from a density p relative to a measure μ on a sample space $(\mathcal{X}, \mathcal{A})$. It is known that p belongs to a collection $\mathcal{P}$ of densities, and the problem is to estimate the value $\chi(p)$ of a functional $\chi\colon \mathcal{P} \to \mathbb{R}$. Our main interest is in the situation of a semiparametric or nonparametric model, where $\mathcal{P}$ is infinite dimensional.

Estimating equations have been found a good strategy for constructing estimators in semiparametric models [2, 34, 40]. Because the model is of (much) higher dimension than the parameter of interest, setting up a good estimating equation often requires an initial estimator $\hat{\eta}_n$ of a "nuisance parameter", and an estimating equation for $\theta = \chi(p)$ may take the form

$$\mathbb{P}_n \psi_{\theta,\hat{\eta}_n} = 0.$$

Here $\mathbb{P}_n f$ is short for $n^{-1}\sum_{i=1}^n f(X_i)$, and $x \mapsto \psi_{\theta,\eta}(x)$ is a given measurable map, for each (θ, η). In the present paper it will be more convenient to work

[†]The research leading to these results has received funding from the European Research Council underERC GrantAgreement320637.

AMS 2000 subject classifications: Primary 62G05, 62G20, 62G20, 62F25

Keywords and phrases: Nonlinear functional, nonparametric estimation, U-statistic, influence function, tangent space

with a one-step version of this estimator, defined by the method of Newton-Rhapson from a linearization of the map $\theta \mapsto \mathbb{P}_n \psi_{\theta,\hat{\eta}_n}$ around an initial estimator $\hat{\theta}_n$, leading to an estimator of the form $\hat{\theta}_n + \hat{V}_n^{-1}\mathbb{P}_n\psi_{\hat{\theta}_n,\hat{\eta}_n}$, for $\hat{V}_n$ an estimate of the derivative of the estimating equation. In more general notation such an estimator can be written as

$$\hat{\chi}_n = \chi(\hat{p}_n) + \mathbb{P}_n \chi_{\hat{p}_n}, \tag{1.1}$$

for $\hat{p}_n$ an initial estimator for p, and $x \mapsto \chi_p(x)$ a given measurable function, for each $p \in \mathcal{P}$.

One possible choice in (1.1) is $\chi_p = 0$, leading to the plug-in estimator $\chi(\hat{p}_n)$. However, unless the initial estimator $\hat{p}_n$ possesses special properties, this choice is typically suboptimal. Better functions χ_p can be constructed by consideration of the *tangent space* of the model. To see this, we write (with $P\chi_{\hat{p}}$ shorthand for $\int \chi_{\hat{p}}(x)\, dP(x)$)

$$\hat{\chi}_n - \chi(p) = \big[\chi(\hat{p}_n) - \chi(p) + P\chi_{\hat{p}_n}\big] + (\mathbb{P}_n - P)\chi_{\hat{p}_n}. \tag{1.2}$$

Because it is properly centered, we may expect the sequence $\sqrt{n}(\mathbb{P}_n - P)\chi_{\hat{p}_n}$ to tend in distribution to a mean-zero normal distribution. The term between square brackets on the right of (1.2), which we shall refer to as the *bias term*, depends on the initial estimator $\hat{p}_n$, and it would be natural to construct the function χ_p such that this term does not contribute to the limit distribution, or at least is not dominating the expression. Thus we would like to choose this function such that the "bias term" is at least of the order $O_P(n^{-1/2})$. A good choice is to ensure that the term $P\chi_{\hat{p}_n}$ acts as minus the first derivative of the functional χ in the "direction" $\hat{p}_n - p$. Functions $x \mapsto \chi_p(x)$ with this property are known as *influence functions* in semiparametric theory [16, 22, 35, 5, 2], go back to the *von Mises calculus* due to [33], and play an important role in robust statistics [14, 11], or [40], Chapter 20.

For an influence function we may expect that the "bias term" is quadratic in the error $d(\hat{p}_n, p)$, for an appropriate distance d. In that case it is certainly negligible as soon as this error is of order $o_P(n^{-1/4})$. Such a "no-bias" condition is well known in semiparametric theory (e.g. condition (25.52) in [40] or (11) in [20]). However, typically it requires that the model $\mathcal{P}$ be "not too big". For instance, a regression or density function on d-dimensional space can be estimated at rate $n^{-1/4}$ if it is a-priori known to have at least $d/2$ derivatives (indeed $\alpha/(2\alpha + d) \geq 1/4$ if $\alpha \geq d/2$). The purpose of this paper is to develop estimation procedures for the case that no estimators exist that attain a $O_P(n^{-1/4})$ rate of convergence. The estimator (1.1) is then suboptimal, because it fails to make a proper trade-off between "bias" and "variance": the two terms in (1.2) have different magnitudes. Our strategy

is to replace the linear term $\mathbb{P}_n\chi_p$ by a general U-statistic $\mathbb{U}_n\chi_p$, for an appropriate m-dimensional *influence function* $(x_1,\ldots,x_m)\mapsto\chi_p(x_1,\ldots,x_m)$, chosen using a von Mises expansion of $p\mapsto\chi(p)$. Here the order m is adapted to the size of the model $\mathcal{P}$ and the type of functional to be estimated.

Unfortunately, "exact" higher-order influence functions turn out to exist only for special functionals χ. To treat general functionals χ we approximate these by simpler functionals, or use approximate influence functions. The rate of the resulting estimator is then determined by a trade-off between bias and variance terms. It may still be of order $1/\sqrt{n}$, but it is often slower. In the former case, surprisingly, one may obtain semiparametric efficiency by estimators whose variance is determined by the linear term, but whose bias is corrected using higher order influence functions.

The conclusion that the "bias term" in (1.2) is quadratic in the estimation error $d(\hat{p}_n,p)$ is based on a worst case analysis. First, there exist a large number of models and functionals of interest that permit a first order influence function that is unbiased in the nuisance parameter. (E.g. adaptive models as considered in [1], models allowing a sufficient statistic for the nuisance parameter as in [38, 39], mixture models as considered in [19, 24, 36], and convex-linear models in survival analysis.) In such models there is no need for higher-order estimating equations. Second, the analysis does not take special, structural properties of the initial estimators $\hat{p}_n$ into account. An alternative approach would be to study the bias of a particular estimator in detail, and adapt the estimating equation to this special estimator. The strategy in this paper is not to use such special properties and focus on estimating equations that work with general initial estimators $\hat{p}_n$.

The motivation for our new estimators stems from studies in epidemiology and econometrics that include covariates whose influence on an outcome of interest cannot be reliably modelled by a simple model. These covariates may themselves not be of interest, but are included in the analysis to adjust the analysis for possible bias. For instance, the mechanism that describes why certain data is missing is in terms of conditional probabilities given several covariates, but the functional form of this dependence is unknown. Or, to permit a causal interpretation in an observational study one conditions on a set of covariates to control for confounding, but the form of the dependence on the confounding variables is unknown. One may hypothesize in such situations that the functional dependence on a set of (continuous) covariates is smooth (e.g. $d/2$ times differentiable in the case of d covariates), or even linear. Then the usual estimators will be accurate (at order $O_P(n^{-1/2})$) if the hypothesis is true, but they will be badly biased in the other case. In particular, the usual normal-theory based confidence intervals may be totally misleading: they will be both too narrow and wrongly located. The

methods in this paper yield estimators with (typically) wider corresponding confidence intervals, but they are correct under weaker assumptions.

The mathematical contributions of the paper are to provide a heuristic for constructing minimax estimators in semiparametric models, and to apply this to a concrete model, which is a template for a number of other models (see [27, 37]). The methods connect to earlier work [13, 21] on the estimation of functionals on nonparametric models, but differs by our focus on functionals that are defined in terms of the structure of a semiparametric model. This requires an analysis of the inverse map from the density of the observations to the parameters, in terms of the semiparametric tangent spaces of the models. Our second order estimators are related to work on quadratic functionals, or functionals that are well approximated by quadratic functionals, as in [10, 15, 3, 4, 17, 18, 6, 7]. While we place the construction of minimax estimators for these special functionals in a wider framework, our focus differs by going beyond quadratic estimators and to consider semiparametric models.

Our mathematical results are in part conditional on a scale of regularity parameters, through a dimension (8.9) and partition of this dimension that depends on two of these parameters. We hope to discuss adaptation to these parameters in future work.

General heuristics of our construction are given in Section 3. Sections 4–8 are devoted to constructing new estimators for the mean response effect in missing data problems. In Section 9 we briefly discuss some other problems, including the problem of estimating a density at a point, where already first order influence functions do not exist and our heuristics naturally lead to projection estimators. Section 10 collects technical proofs. Sections 11, 12 and 13 (in the supplement [26]) discuss three key concepts of the paper: influence functions, projections and U-statistics.

2. Notation. Let $\mathbb{U}_n$ denote the *empirical U-statistic* measure, viewed as an operator on functions. For given $k \leq n$ and a function $f: \mathcal{X}^k \to \mathbb{R}$ on the sample space this is defined by

$$\mathbb{U}_n f = \frac{1}{n(n-1)\cdots(n-k+1)} \sum\sum\cdots\sum_{1\leq i_1\neq i_2\neq\cdots\neq i_k\leq n} f(X_{i_1}, X_{i_2}, \cdots, X_{i_k}).$$

We do not let the order k show up in the notation $\mathbb{U}_n f$. This is unnecessary, as the notation is consistent in the following sense: if a function $f: \mathcal{X}^l \to \mathbb{R}$ of $l < k$ arguments is considered a function of k arguments that is constant in its last $k-l$ arguments, then the right side of the preceding display is well defined and is exactly the corresponding U-statistic of order l. In particular,

$\mathbb{U}_n f$ is the *empirical distribution* $\mathbb{P}_n$ applied to f if $f\colon \mathcal{X} \to \mathbb{R}$ depends on only one argument.

We write $P^n \mathbb{U}_n f = P^k f$ for the expectation of $\mathbb{U}_n f$ if $X_1, \ldots, X_n$ are distributed according to the probability measure P. We also use this operator notation for the expectations of statistics in general.

We call f *degenerate* relative to P if $\int f(x_1, \ldots, x_k)\, dP(x_i) = 0$ for every i and every $(x_j\colon j \neq i)$, and we call f *symmetric* if $f(x_1, \ldots, x_k)$ is invariant under permutation of the arguments $x_1, \ldots, x_k$. Given an arbitrary measurable function $f\colon \mathcal{X}^k \to \mathbb{R}$ we can form a function that is degenerate relative to P by subtracting the orthogonal projection in $L_2(P^k)$ onto the functions of at most $k-1$ variables. This degenerate function can be written in the form (e.g. [40], Lemma 11.11)

$$(D_P f)(X_1, \ldots, X_k) = \sum_{A \subset \{1,\ldots,k\}} (-1)^{k-|A|} \mathrm{E}_P\Big[f(X_1, \ldots, X_k) \big|\, X_i\colon i \in A \Big], \tag{2.1}$$

where the sum if over all subsets A of $\{1, \ldots, k\}$, including the empty set, for which the conditional expectation is understood to be $P^k f$. If the function f is symmetric, then so is the function $D_P f$.

Given two functions $g, h\colon \mathcal{X} \to \mathbb{R}$ we write $g \times h$ for the function $(x, y) \mapsto g(x)h(y)$. More generally, given k functions $g_1, \ldots, g_k$ we write $g_1 \times \cdots \times g_k$ for the tensor product of these functions. Such product functions are degenerate iff all functions in the product have mean zero.

A *kernel operator* $K\colon L_r(\mathcal{X}, \mathcal{A}, \mu) \to L_r(\mathcal{X}, \mathcal{A}, \mu)$ takes the form $(Kf)(x) = \int \bar{K}(x, y) f(y)\, d\mu(y)$ for some measurable function $\bar{K}\colon \mathcal{X}^2 \to \mathbb{R}$. We shall abuse notation in denoting the operator K and the *kernel* $\bar{K}$ with the same symbol: $K = \bar{K}$. A (weighted) projection on a finite-dimensional space is a kernel operator. We discuss such projections in Section 12.

The set of measurable functions whose rth absolute power is μ-integrable is denoted $L_r(\mu)$, with norm $\|\cdot\|_{r,\mu}$, or $\|\cdot\|_r$ if the measure is clear; or also as $L_r(w)$ with norm $\|\cdot\|_{r,w}$ if w is a density relative to a given dominating measure. For $r = \infty$ the notation $\|\cdot\|_\infty$ refers to the uniform norm.

3. Heuristics. Our basic estimator has the form (1.1) except that we replace the linear term by a general U-statistic. Given measurable functions $\chi_p\colon \mathcal{X}^m \to \mathbb{R}$, for a fixed *order* m, we consider estimators $\hat{\chi}_n$ of $\chi(p)$ of the type

$$\hat{\chi}_n = \chi(\hat{p}_n) + \mathbb{U}_n \chi_{\hat{p}_n}. \tag{3.1}$$

The initial estimators $\hat{p}_n$ are thought to have a certain (optimal) convergence rate $d(\hat{p}_n, p) \to 0$, but need not possess (further) special properties.

Throughout we shall treat these estimators as being based on an independent sample of observations, so that the stochasticities in (3.1) present in $\hat{p}_n$ and $\mathbb{U}_n$ are independent. This takes away technical complications, and allows us to focus on rates of estimation in full generality. (A simple way to avoid the resulting asymmetry would be to swap the two samples, calculate the estimator a second time and take the average.)

3.1. *Influence functions.* The key is to find suitable "influence functions" χ_p. A decomposition of type (1.2) for the estimator (3.1) yields

$$\hat{\chi}_n - \chi(p) = \left[\chi(\hat{p}_n) - \chi(p) + P^m \chi_{\hat{p}_n}\right] + (\mathbb{U}_n - P^n)\chi_{\hat{p}_n}. \tag{3.2}$$

This suggests to construct the influence functions such that $-P^m \chi_{\hat{p}_n}$ represents the first m terms of the Taylor expansion of $\chi(\hat{p}_n) - \chi(p)$. First this implies that the influence function used in (3.1) must be unbiased:

$$P^m \chi_p = 0.$$

Next, to operationalize a "Taylor expansion" on the (infinite-dimensional) "manifold" $\mathcal{P}$ we employ "smooth" submodels $t \mapsto p_t$ mapping a neighbourhood of $0 \in \mathbb{R}$ to $\mathcal{P}$ and passing through p at $t = 0$ (i.e. $p_0 = p$). We determine χ_p such that for each chosen submodel

$$\frac{d^j}{dt^j}_{|t=0} \chi(p_t) = -\frac{d^j}{dt^j}_{|t=0} P^m \chi_{p_t}, \qquad j = 1, \ldots, m.$$

A slight strengthening is to impose this condition "everywhere" on the path, i.e. the jth derivative of $t \mapsto \chi(p_t)$ at t is the jth derivative of $h \mapsto -P_t^m \chi_{p_{t+h}}$ at $h = 0$, for every t. If the map $(s,t) \mapsto P_s^m \chi_{p_t}$ is smooth, then the latter implies (cf. Lemma 11.1 applied with $\chi = f$ and $g(s,t) = -P_t^m \chi_{p_s}$)

$$\frac{d^j}{dt^j}_{|t=0} \chi(p_t) = \frac{d^j}{dt^j}_{|t=0} P_t^m \chi_p, \qquad j = 1, \ldots, m. \tag{3.3}$$

Relative to the previous formula the subscript t on the right hand side has changed places, and the negative sign has disappeared. Under regularity conditions equation (3.3) for $m = 1$ can be written in the form

$$\frac{d}{dt}_{|t=0} \chi(p_t) = \frac{d}{dt}_{|t=0} P_t \chi_p = P \chi_p g, \tag{3.4}$$

where $g = (d/dt)_{t=0} p_t / p$ is the *score function* of the model $t \mapsto p_t$ at $t = 0$.

A function χ_p satisfying (3.3) is exactly what is called an *influence function* in semiparametric theory: a function in $L_2(p)$ whose inner products with

the elements of the tangent space (the set of score functions g of appropriate submodels $t \mapsto p_{g,t}$) represent the derivative of the functional ([40], page 363, or [2, 22, 39, 35]). For $m > 1$ equation (3.3) can be expanded similarly in terms of inner products of the influence function with score functions, but higher-order score functions arise next to ordinary score functions. Suitable higher-order tangent spaces are discussed in [27] (also see [37]), using score functions as defined in [43]. A discussion of *second order* scores and tangent spaces can be found in [28]. Second order tangent spaces are also discussed in [23], from a different point of view and with a different purpose.

Here we take a different route, defining higher order influence functions as influence functions of lower order ones. Any mth order, zero-mean U-statistic can be decomposed as the sum of m degenerate U-statistics of orders $1, 2, \ldots, m$, by way of its Hoeffding decomposition. In the present situation we can write

$$\mathbb{U}_n \chi_p = \mathbb{U}_n \chi_p^{(1)} + \tfrac{1}{2} \mathbb{U}_n \chi_p^{(2)} + \cdots + \frac{1}{m!} \mathbb{U}_n \chi_p^{(m)},$$

where $\chi_p^{(j)} \colon \mathcal{X}^j \to \mathbb{R}$ is a degenerate kernel of j arguments, defined uniquely as a projection of χ_p (cf. [42] and (2.1)). (At $m = n$ the left side evaluates to the symmetrized function χ_p, which is expressed in the functions $\chi_p^{(j)}$.) Suitable functions $\chi_p^{(j)}$ in this decomposition can be found by the following algorithm:

[1] Let $x_1 \mapsto \bar{\chi}_p^{(1)}(x_1)$ be a first order influence function of the functional $p \mapsto \chi(p)$.
[2] Let $x_j \mapsto \bar{\chi}_p^{(j)}(x_1, \ldots, x_j)$ be a first order influence function of the functional $p \mapsto \bar{\chi}_p^{(j-1)}(x_1, \ldots, x_{j-1})$, for each $x_1, \ldots, x_{j-1}$, and $j = 2, \ldots, m$.
[3] Let $\chi_p^{(j)} = D_P \bar{\chi}_p^{(j)}$ be the degenerate part of $\bar{\chi}_p^{(j)}$ relative to P, as defined in (2.1).

See Lemma 11.2 for a proof. Thus higher order influence functions are constructed as first order influence functions of influence functions. Somewhat abusing language we shall refer to the function $\chi_p^{(j)}$ also as a "jth order influence function". The overall order m will be fixed at a suitable value; for simplicity we do not let this show up in the notation χ_p.

Because only inner products with scores matter, an "influence function" is unique only up to projections onto the tangent space and (averaging over) permutations of its arguments. With any choice of influence functions the algorithm produces some influence function. In particular, the starting influence function $\bar{\chi}_p^{(1)}$ in step [1] may be any function χ_p that satisfies (3.4)

for every score g; it does not have to possess mean zero, or be an element of the tangent space. A similar remark applies to the (first order) influence functions found in step [2]. It is only in step [3] that we make the influence functions degenerate.

3.2. *Bias-variance trade-off.* Because it is centered, the "variance part" in (3.2), the variable $(\mathbb{U}_n - P^n)\chi_{\hat{p}_n}$, should not change noticeably if we replace $\hat{p}_n$ by p, and be of the same order as $(\mathbb{U}_n - P^n)\chi_p$. For a fixed square-integrable function χ_p the latter centered U-statistic is well known to be of order $O_P(n^{-1/2})$, and asymptotically normal if suitably scaled. A completely successful representation of the "bias" $R_n = \chi(\hat{p}_n) - \chi(p) + P^m\chi_{\hat{p}_n}$ in (3.2) would lead to an error $R_n = O_P\big(d(\hat{p}_n, p)^{m+1}\big)$, which becomes smaller with increasing order m. Were this achievable for any m, then a $\sqrt{n}$-estimator would exist no matter how slow the convergence rate $d(\hat{p}_n, p)$ of the initial estimator. Not surprisingly, in many cases of interest this ideal situation is not real. This is due to the non-existence of influence functions that can exactly represent the Taylor expansion of $\chi(\hat{p}_n) - \chi(p)$. The technical reason is that multi-linear maps (even smooth ones) may not be representable through kernels.

In general, we have to content ourselves with a partial representation. Next to a remainder term of order $O_P\big(d(\hat{p}_n, p)^{m+1}\big)$, we then incur a "representation bias". The latter bias can be made arbitrarily small by choice of the influence function, but only at the cost of increasing its variance. We thus obtain a trade-off between a variance and two biases. This typically results in a variance that is larger than $1/n$, and a rate of convergence that is slower than $1/\sqrt{n}$, although sometimes a nontrivial bias correction is possible without increasing the variance.

3.3. *Approximate functionals.* An attractive method to find approximating influence functions is to compute exact influence functions for an approximate functional. Because smooth functionals on finite-dimensional models typically possess influence functions to any order, projections on finite-dimensional models may deliver such approximations.

A simple approximation would be $\chi(\widetilde{p})$ for a given map $p \mapsto \widetilde{p}$ mapping the model $\mathcal{P}$ onto a suitable "smaller" model $\widetilde{\mathcal{P}}$ (typically a submodel $\widetilde{\mathcal{P}} \subset \mathcal{P}$). A closer approximation can be obtained by also including a derivative term. Consider the functional $\widetilde{\chi}\colon \mathcal{P} \to \mathbb{R}$ defined by, for a given map $p \mapsto \widetilde{p}$,

$$\widetilde{\chi}(p) = \chi(\widetilde{p}) + P\chi^{(1)}_{\widetilde{p}}. \tag{3.5}$$

(A more complete notation would be $\widetilde{p}(p)$; the right hand side depends on p in three ways.) By the definition of an influence function the term $-P\chi^{(1)}_{\widetilde{p}}$

acts as the first order Taylor expansion of $\chi(\widetilde{p}) - \chi(p)$. Consequently, we may expect that

$$\bigl|\widetilde{\chi}(p) - \chi(p)\bigr| = O\bigl(d(\widetilde{p}, p)^2\bigr). \tag{3.6}$$

This ought to be true for any "projection" $p \mapsto \widetilde{p}$. If we choose the projection such that, for any path $t \mapsto p_t$,

$$\frac{d}{dt}_{|t=0}\Bigl(\chi(\widetilde{p}_t) + P_0\chi^{(1)}_{\widetilde{p}_t}\Bigr) = 0, \tag{3.7}$$

then the functional $p \mapsto \widetilde{\chi}(p)$ will be locally (around p_0) equivalent to the functional $p \mapsto \chi\bigl(\widetilde{p}_0\bigr) + P\chi^{(1)}_{\widetilde{p}_0}$ (which depends on p in only one place, p_0 being fixed) in the sense that the first order influence functions are the same. The first order influence function of the second, linear functional at p_0 is equal to $\chi^{(1)}_{\widetilde{p}_0}$, and hence for a projection satisfying (3.7) the first order influence function of the functional $p \mapsto \widetilde{\chi}(p)$ will be

$$\widetilde{\chi}^{(1)}_p = \chi^{(1)}_{\widetilde{p}}. \tag{3.8}$$

In words, this means that the influence function of the approximating functional $\widetilde{\chi}$ satisfying (3.5) and (3.7) at p is obtained by substituting $\widetilde{p}$ for p in the influence function of the original functional.

This is relevant when obtaining higher order influence functions. As these are recursive derivatives of the first order influence function (see [1]–[3] in Section 3.1), the preceding display shows that we must compute influence functions of

$$p \mapsto \chi^{(1)}_{\widetilde{p}}(x),$$

i.e. we "differentiate on the model $\widetilde{\mathcal{P}}$". If the latter model is sufficiently simple, for instance finite-dimensional, then exact higher order influence functions of the functional $p \mapsto \widetilde{\chi}(p)$ ought to exist. We can use these as approximate influence functions of $p \mapsto \chi(p)$.

4. Estimating the mean response in missing data models. Suppose that a typical observation is distributed as $X = (YA, A, Z)$, for Y and A taking values in the two-point set $\{0, 1\}$ and conditionally independent given Z.

This model is standard in biostatistical applications, with Y an "outcome" or "response variable", which is observed only if the indicator A takes the value 1. The covariate Z is chosen such that it contains all information on the dependence between the response and the missingness indicator A,

thus making the response *missing at random*. Alternatively, we think of Y as a "counterfactual" outcome if a treatment were given ($A = 1$) and estimate (half) the treatment effect under the assumption of *no unmeasured confounders*.

The model can be parameterized by the marginal density f of Z (relative to some dominating measure ν) and the probabilities $b(z) = \mathrm{P}(Y = 1 | Z = z)$ and $a(z)^{-1} = \mathrm{P}(A = 1 | Z = z)$. (Using a for the inverse probability simplifies later formulas.) Alternatively, the model can be parameterized by the pair (a, b) and the function $g = f/a$, which is the conditional density of Z given $A = 1$, up to the norming factor $\mathrm{P}(A = 1)$. Thus the density p of an observation X is described by the triplet (a, b, f), or equivalently the triplet (a, b, g). For simplicity of notation we write p instead of $p_{a,b,f}$ or $p_{a,b,g}$, with the implicit understanding that a generic p corresponds one-to-one to a generic (a, b, f) or (a, b, g).

We wish to estimate the *mean response* $\mathrm{E}Y = \mathrm{E}b(Z)$, i.e. the functional

$$\chi(p) = \int bf \, d\nu = \int abg \, d\nu.$$

Estimators that are $\sqrt{n}$-consistent and asymptotically efficient in the semiparametric sense have been constructed using a variety of methods (e.g. [30, 31]), but only if a or b, or both, parameters are restricted to sufficiently small regularity classes. For instance, if the covariate Z ranges over a compact, convex subset $\mathcal{Z}$ of $\mathbb{R}^d$, then the mentioned papers provide $\sqrt{n}$-consistent estimators under the assumption that a and b belong to Hölder classes $C^\alpha(\mathcal{Z})$ and $C^\beta(\mathcal{Z})$ with α and β large enough that

$$\frac{\alpha}{2\alpha + d} + \frac{\beta}{2\beta + d} \geq \frac{1}{2}. \tag{4.1}$$

(See e.g. Section 2.7.1 in [41] for the definition of Hölder classes). For moderate to large dimensions d this is a restrictive requirement. In the sequel we consider estimation for arbitrarily small α and β.

Throughout we assume that the parameters a, b and g are contained in Hölder spaces $C^\alpha(\mathcal{Z})$, $C^\beta(\mathcal{Z})$ and $C^\gamma(\mathcal{Z})$ of functions on a compact, convex domain in $\mathbb{R}^d$. We derive two types of results:

(a) In Section 7 we show that a $\sqrt{n}$-rate is attainable by using a higher order estimating equation (of order determined by γ) as long as

$$\frac{\alpha + \beta}{2} \geq \frac{d}{4}. \tag{4.2}$$

This condition is strictly weaker than the condition (4.1) under which the linear estimator attains a $\sqrt{n}$-rate. Thus even in the $\sqrt{n}$-situation

higher order estimating equations may yield estimators that are applicable in a wider range of models. For instance, in the case that $\alpha = \beta$ the cut-off (4.1) arises for $\alpha = \beta \geq d/4$, whereas (4.2) reduces to $\alpha = \beta \geq d/2$.

(b) We consider minimax estimation in the case $(\alpha + \beta)/2 < d/4$, when the rate becomes slower than $1/\sqrt{n}$. It is shown in [29] that even if $g = f/a$ were known, then the minimax rate for a and b ranging over balls in the Hölder classes $C^\alpha(\mathcal{Z})$ and $C^\beta(\mathcal{Z})$ cannot be faster than $n^{-(2\alpha+2\beta)/(2\alpha+2\beta+d)}$. In Section 8 we show that this rate is attainable if g is known, and also if g is unknown, but is a-priori known to belong to a Hölder class $C^\gamma(\mathcal{Z})$ for sufficiently large γ, as given by (8.11). (Heuristic arguments, not discussed in this paper, appear to indicate that for smaller γ the minimax rate is slower than $n^{-(2\alpha+2\beta)/(2\alpha+2\beta+d)}$.)

After reviewing the tangent space and first order theory in Section 4.1, we discuss the second order estimator separately in Section 5. The preceding results are next obtained in Sections 7 ($\sqrt{n}$-rate if $(\alpha + \beta)/2 \geq d/4$) and 8 (slower rate if $(\alpha + \beta)/2 < d/4$), using the higher-order influence functions of an approximate functional, which is defined in the intermediate Section 6.

ASSUMPTION 4.1. *We assume throughout that the functions* $1/a, b, g$ *and their preliminary estimators* $1/\hat{a}, \hat{b}, \hat{g}$ *are bounded away from their extremes: 0 and 1 for the first two, and 0 and* ∞ *for the third.*

4.1. *Tangent space and first order influence function.* The one-dimensional submodels $t \mapsto p_t$ induced by paths of the form $a_t = a + t\alpha$, $b_t = b + t\beta$, and $f_t = f(1 + t\phi)$ for given directions α, β and ϕ (where $\int \phi f \, d\nu = 0$) yield score functions

$$\begin{aligned} B_p^a \alpha(X) &= -\frac{Aa(Z) - 1}{a(Z)(a-1)(Z)} \alpha(Z), && a - \text{score}, \\ B_p^b \beta(X) &= \frac{A\big(Y - b(Z)\big)}{b(Z)(1-b)(Z)} \beta(Z), && b - \text{score}, \\ B_p^f \phi(X) &= \phi(Z), && f - \text{score}. \end{aligned}$$

Here B_p^a, B_p^b, B_p^f are the *score operators* for the three parameters, whose direct sum is the overall score operator, which we write as B_p: $B_p(\alpha, \beta, \phi)(X)$ is the sum of the three left sides of the preceding equation. The first-order influence function is well known to take the form

$$\chi_p^{(1)}(X) = Aa(Z)\big(Y - b(Z)\big) + b(Z) - \chi(p). \tag{4.3}$$

Indeed, it is straightforward to verify that this function satisfies, for every path $t \mapsto p_t$ as described previously,

$$\frac{d}{dt}_{|t=0} \chi(p_t) = \mathrm{E}_p \chi_p^{(1)}(X)\, B_p(\alpha, \beta, \phi)(X).$$

The advantage of choosing a an inverse probability is clear from the form of the (random part of the) influence function, which is bilinear in (a, b). The corresponding "first order bias" can be computed to be

$$\chi(\hat{p}) - \chi(p) + P\chi_{\hat{p}}^{(1)} = -\int (\hat{a} - a)(\hat{b} - b)\, g\, d\nu. \tag{4.4}$$

In agreement with the heuristics given in Sections 1 and 3 this bias is quadratic in the errors of the initial estimator.

Actually, the form of the bias term is special in that square estimation errors $(\hat{a} - a)^2$ and $(\hat{b} - b)^2$ of the two initial estimators $\hat{a}$ and $\hat{b}$ do no arise, but only the product $(\hat{a} - a)(\hat{b} - b)$ of their errors. This property, termed "double robustness" in [34], makes that for first order inference it suffices that one of the two parameters be estimated well. A prior assumption that the parameters a and b are α and β regular, respectively, would allow estimation errors of the orders $n^{-\alpha/(2\alpha+d)}$ and $n^{-\beta/(2\beta+d)}$. If the product of these rates is $O(n^{-1/2})$, then the bias term matches the variance. This leads to the (unnecessarily restrictive) condition (4.1).

If the preliminary estimators $\hat{a}$ and $\hat{b}$ are solely selected for having small errors $\|\hat{a} - a\|$ and $\|\hat{b} - b\|$ (e.g. minimax in the L_2-norm), then it is hard to see why (4.4) would be small unless the product $\|\hat{a} - a\|\|\hat{b} - b\|$ of the errors is small. Special estimators might exploit that the bias is an integral, in which cancellation of errors could occur. As we do not wish to use special estimators, our approach will be to replace the linear estimating equation by a higher order one, leading to an analogue of (4.4) that is a cubic or higher order polynomial of the estimation errors.

It may be noted that the marginal density f (or g) does not enter into the first order influence function (4.3). Even though the functional depends on f (or g), a rate on the initial estimator of this function is not needed for the construction of the first order estimator. This will be different at higher orders.

5. Second order estimator. In this section we derive a second order influence function for the missing data problem, and analyze the risk of the corresponding estimator. This estimator is minimax if $(\alpha + \beta)/2 \geq d/4$ and

$$\frac{\gamma}{2\gamma + d} \geq \frac{1}{2} \wedge \frac{2\alpha + 2\beta}{d + 2\alpha + 2\beta} - \frac{\alpha}{2\alpha + d} - \frac{\beta}{2\beta + d}. \tag{5.1}$$

In the other case, higher order estimators have smaller risk, as shown in Sections 7-8. However, it is worth while to treat the second order estimator separately, as its construction exemplifies essential elements, without involving technicalities attached to the higher order estimators.

To find a second order influence function, we follow the strategy [1]–[3] of Section 3.1, and try and find a function $\chi_p^{(2)}\colon \mathcal{X}^2 \to \mathbb{R}$ such that, for every $x_1 = (y_1 a_1, a_1, z_1)$, and all directions α, β, ϕ,

$$\frac{d}{dt}_{|t=0}\Big[\chi_{p_t}^{(1)}(x_1) + \chi(p_t)\Big] = \mathrm{E}_p \chi_p^{(2)}(x_1, X_2)\, B_p(\alpha,\beta,\phi)(X_2).$$

Here the expectation E_p on the right side is relative to the variable X_2 only, with x_1 fixed. This equation expresses that $x_2 \mapsto \chi_p^{(2)}(x_1, x_2)$ is a first order influence function of $p \mapsto \chi_p^{(1)}(x_1) + \chi(p)$, for fixed x_1. On the left side we added the "constant" $\chi(p_t)$ to the first order influence function (giving another first order influence function) to facilitate the computations. This is justified as the strategy [1]–[3] works with any influence function. In view of (4.3) and the definitions of the paths $t \mapsto a+t\alpha$, $t \mapsto b+t\beta$ and $t \mapsto f(1+t\phi)$, this leads to the equation

$$\begin{aligned} a_1\big(y_1 - b(z_1)\big)\alpha(z_1) - \big(a_1 a(z_1) - 1\big)\beta(z_1) & \\ = \mathrm{E}_p \chi_p^{(2)}(x_1, X_2)\, B_p(\alpha,\beta,\phi)(X_2). & \end{aligned} \tag{5.2}$$

Unfortunately, no function $\chi_p^{(2)}$ that solves this equation for every (α, β, ϕ) exists. To see this note that for the special triplets with $\beta = \phi = 0$ the requirement can be written in the form

$$\alpha(z_1) = \mathrm{E}_p\left[\frac{\chi_p^{(2)}(x_1, X_2)}{a_1\big(y_1 - b(z_1)\big)} \frac{1 - A_2 a(Z_2)}{a(Z_2)(a-1)(Z_2)}\right] \alpha(Z_2).$$

The right side of the equation can be written as $\int K(z_1, z_2)\alpha(z_2)\, dF(z_2)$, for $K(z_1, Z_2)$ the conditional expectation of the function in square brackets given Z_2. Thus it is the image of α under the kernel operator with kernel K. If the equation were true for any α, then this kernel operator would work as the identity operator. However, on infinite-dimensional domains the identity operator is not given by a kernel. (Its kernel would be a "Dirac function on the diagonal".)

Therefore, we have to be satisfied with an influence function that gives a partial representation only. In particular, a projection onto a finite-dimensional linear space possesses a kernel, and acts as the identity on this linear space. A "large" linear space gives representation in "many"

directions. By reducing the expectation in (5.2) to an integral relative to the marginal distribution of Z_2, we can use an orthogonal projection $\Pi_p\colon L_2(g) \to L_2(g)$ onto a subspace L of $L_2(g)$. Writing also Π_p for its kernel, and letting $S_2 h$ denote the symmetrization $\big(h(X_1, X_2) + h(X_2, X_1)\big)/2$ of a function $h\colon \mathcal{X}^2 \to \mathbb{R}$, we define

$$\chi_p^{(2)}(X_1, X_2) = -2S_2\Big[A_1\big(Y_1 - b(Z_1)\big)\Pi_p(Z_1, Z_2)\big(A_2 a(Z_2) - 1\big)\Big]. \tag{5.3}$$

LEMMA 5.1. *For $\chi_p^{(2)}$ defined by (5.3) with Π_p the kernel of an orthogonal projection $\Pi_p\colon L_2(g) \to L_2(g)$ onto a subspace $L \subset L_2(g)$, equation (5.2) is satisfied for every path $t \mapsto p_t$ corresponding to directions (α, β, ϕ) such that $\alpha \in L$ and $\beta \in L$.*

PROOF. By definition $\mathrm{E}(A|\, Z) = (1/a)(Z)$ and $\mathrm{E}(Y|\, Z) = b(Z)$. Also $\mathrm{var}\big(Aa(Z)|\, Z\big) = a(Z) - 1$ and $\mathrm{var}(Y|\, Z) = b(Z)(1 - b)(Z)$. By direct computation using these identities, we find that for the influence function (5.3) the right side of (5.2) reduces to

$$a_1\big(y_1 - b(z_1)\big)\Pi_p\alpha(z_1) - \big(a_1 a(z_1) - 1\big)\Pi_p\beta(z_1).$$

Thus (5.2) holds for every (α, β, ϕ) such that $\Pi_p\alpha = \alpha$ and $\Pi_p\beta = \beta$. ∎

Together with the first order influence function (4.3) the influence function (5.3) defines the (approximate) influence function $\chi_p = \chi_p^{(1)} + \frac{1}{2}\chi_p^{(2)}$. For an initial estimator $\hat{p}$ based on independent observations we now construct the estimator (3.1), i.e.

$$\hat{\chi}_n = \chi(\hat{p}) + \mathbb{P}_n\chi_{\hat{p}}^{(1)} + \tfrac{1}{2}\mathbb{U}_n\chi_{\hat{p}}^{(2)}. \tag{5.4}$$

Unlike the first order influence function, the second order influence function does depend on the density f of the covariates, or rather the function $g = f/a$ (through the kernel Π_p, which is defined relative to $L_2(g)$), and hence the estimator (5.4) involves a preliminary estimator of g. As a consequence, the quality of the estimator of the functional χ depends on the precision by which g (as part of the plug-in $\hat{p} = (\hat{a}, \hat{b}, \hat{g})$) can be estimated.

Let $\hat{\mathrm{E}}_p$ and $\hat{\mathrm{var}}_p$ denote conditional expectations given the observations used to construct $\hat{p}$, let $\|\cdot\|_r$ be the norm of $L_r(g)$, and let $\|\Pi\|_r$ denote the norm of an operator $\Pi\colon L_r(g) \to L_r(g)$.

THEOREM 5.1. *The estimator $\hat{\chi}_n$ given in (5.4) with influence functions $\chi_p^{(1)}$ and $\chi_p^{(2)}$ defined by (4.3) and (5.3), for Π_p the kernel of an orthogonal*

projection in $L_2(g)$ onto a k-dimensional linear subspace, satisfies, for $r \geq 2$ (with $r/(r-2) = \infty$ if $r = 2$),

$$\hat{\mathrm{E}}_p\hat{\chi}_n - \chi(p) = O_P\Big(\|\Pi_p\|_r\|\Pi_{\hat{p}}\|_r\|\hat{a}-a\|_r\|\hat{b}-b\|_r\|\hat{g}-g\|_{r/(r-2)}\Big) + O_P\Big(\big\|(I-\Pi_p)(a-\hat{a})\big\|_2\big\|(I-\Pi_p)(b-\hat{b})\big\|_2\Big),$$

$$\hat{\mathrm{var}}_p\hat{\chi}_n = O_P\Big(\frac{1}{n} + \frac{k}{n^2}\Big).$$

The two terms in the bias result from having to estimate p in the second order influence function (giving "third order bias") and using an approximate influence function (leaving the remainders $I - \Pi_p$ after projection), respectively. The terms $1/n$ and k/n^2 in the variance appear as the variances of $\mathbb{U}_n\chi_p^{(1)}$ and $\mathbb{U}_n\chi_p^{(2)}$, the second being a degenerate second order U-statistic (giving $1/n^2$, see (13.1)) with a kernel of variance k.

The proof of the theorem is deferred to Section 10.1.

Assume now that the range space of the projections Π_p can be chosen such that, for some constant C,

$$\|a - \Pi_p a\|_2 \leq C\Big(\frac{1}{k}\Big)^{\alpha/d}, \qquad \|b - \Pi_p b\|_2 \leq C\Big(\frac{1}{k}\Big)^{\beta/d}. \tag{5.5}$$

Furthermore, assume that there exist estimators $\hat{a}$ and $\hat{b}$ and $\hat{g}$ that achieve convergence rates $n^{-\alpha/(2\alpha+d)}$, $n^{-\beta/(2\beta+d)}$ and $n^{-\gamma/(2\gamma+d)}$, respectively, in $L_r(g)$ and $L_{r/(r-2)}(g)$, uniformly over these a-priori models and a model for g (e.g. for $r = 3$), and that the preceding displays also hold for $\hat{a}$ and $\hat{b}$. These assumptions are satisfied if the unknown functions a and b are "regular" of orders α and β on a compact subset of $\mathbb{R}^d$ (see e.g. [32]). Then the estimator $\hat{\chi}_n$ of Theorem 5.1 attains the square rate of convergence

$$\Big(\frac{1}{n}\Big)^{2\alpha/(2\alpha+d)+2\beta/(2\beta+d)+2\gamma/(2\gamma+d)} \vee \Big(\frac{1}{k}\Big)^{(2\alpha+2\beta)/d} \vee \frac{1}{n} \vee \frac{k}{n^2}. \tag{5.6}$$

We shall see in the next section that the first of the four terms in this maximum can be made smaller by choosing an estimating equation of order higher than 2, while the other three terms arise at any order. This motivates to determine a "second order 'optimal" value of k by balancing the second, third and fourth terms. We next would use the second order estimator if γ is large enough so that the first term is negligible relative to the other terms.

For $(\alpha+\beta)/2 \geq d/4$ we can choose $k = n$ and the resulting rate (the square root of (5.6)) is $n^{-1/2}$ provided that (5.1) holds. The latter condition is certainly satisfied under the sufficient condition (4.1) for the linear estimator to yield rate $n^{-1/2}$.

More interestingly, for $(\alpha+\beta)/2 < d/4$ we choose $k \sim n^{2d/(d+2\alpha+2\beta)}$ and obtain the rate, provided that (5.1) holds,

$$n^{-(2\alpha+2\beta)/(d+2\alpha+2\beta)}.$$

This rate is slower than $n^{-1/2}$, but better than the rate $n^{-\alpha/(2\alpha+d)-\beta/(2\beta+d)}$ obtained by the linear estimator. In [29] this rate is shown to be the fastest possible in the minimax sense, for the model in which a and b range over balls in $C^\alpha(\mathcal{Z})$ and $C^\beta(\mathcal{Z})$, and g being known.

In both cases the second order estimator is better than the linear estimator, but minimax only for sufficiently large γ. This motivates to consider higher order estimators.

6. Approximate functional. Even though the functional of interest does not possess an exact second-order influence function, we might proceed to higher orders by differentiating the approximate second-order influence function $\chi_p^{(2)}$ given in (5.3), and balancing the various terms obtained. However, the formulas are much more transparent if we compute *exact* higher-order influence functions of an approximating functional instead. In this section we first define a suitable functional and next compute its influence functions.

Following the heuristics of Section 3.3, we define an approximate functional by equation (3.5), using a particular projection $p \mapsto \widetilde{p}$ of the parameters. We choose this projection to map the parameters a and b onto finite-dimensional models and leave the parameter g unaltered: p is mapped into an element $\widetilde{p}$ of the approximating model, or equivalently a triplet (a,b,g) into a triplet $(\widetilde{a},\widetilde{b},g)$ in the approximating model for the three parameters (where g is unaltered). (Even though this is not evident in the notation, the projection is joint in the three parameters: the induced maps $(a,b,g) \mapsto \widetilde{a}$ and $(a,b,g) \mapsto \widetilde{b}$ do not reduce to maps $a \mapsto \widetilde{a}$ and $b \mapsto \widetilde{b}$, but $\widetilde{a}$ and $\widetilde{b}$ depend on the full triplet (a,b,g).)

As "model" for (a,b) we consider the product of two affine linear spaces

$$\big(\hat{a} + \underline{a}L\big) \times \big(\hat{b} + \underline{b}L\big), \tag{6.1}$$

for a given finite-dimensional subspace L of $L_2(\nu)$ and fixed functions $\hat{a}, \underline{a}, \hat{b}, \underline{b}\colon \mathcal{Z} \to \mathbb{R}$ that are bounded away from zero and infinity. (Later the functions $\hat{a}$ and $\hat{b}$ are taken equal to the preliminary estimators; one choice for the other functions is $\underline{a} = \underline{b} = 1$.) The pair $(\widetilde{a},\widetilde{b})$ of projections are defined as elements of the model (6.1) satisfying equation (3.7). In view of (4.4), for any path $\widetilde{p}_t \leftrightarrow (\widetilde{a}_t, \widetilde{b}_t, g) = (\widetilde{a} + t\underline{a}\,l, \widetilde{b} + t\underline{b}\,l', g)$, for given $l, l' \in L$,

$$\chi(\widetilde{p}_t) + P\chi^{(1)}_{\widetilde{p}_t} = \chi(p) - \int \big(\widetilde{a} + t\underline{a}\,l - a\big)\big(\widetilde{b} + t\underline{b}\,l' - b\big)\, g\, d\nu. \tag{6.2}$$

Equation (3.7) requires that the derivative of this expression with respect to t at $t = 0$ vanishes. Thus the functions $\widetilde{a}$ and $\widetilde{b}$ must be chosen to satisfy the set of stationary equations, for every $l, l' \in L$,

$$0 = \int (\widetilde{a} - a)\underline{b}\, l'\, g\, d\nu = \int \Big(\frac{\widetilde{a} - \hat{a}}{\underline{a}} - \frac{a - \hat{a}}{\underline{a}}\Big)\, l'\, \underline{abg}\, d\nu, \qquad l' \in L, \tag{6.3}$$

$$0 = \int \underline{a}\, l(\widetilde{b} - b)\, g\, d\nu = \int \Big(\frac{\widetilde{b} - \hat{b}}{\underline{b}} - \frac{b - \hat{b}}{\underline{b}}\Big)\, l\, \underline{abg}\, d\nu, \qquad l \in L. \tag{6.4}$$

Because the functions $(\widetilde{a} - \hat{a})/\underline{a}$ and $(\widetilde{b} - \hat{b})/\underline{b}$ are required to be in L, the second way of writing these equations shows that the latter two functions are the orthogonal projections of the functions $(a - \hat{a})/\underline{a}$ and $(b - \hat{b})/\underline{b}$ onto L in $L_2(\underline{abg})$.

As explained in Section 3.3, as it satisfies (3.7) the projection $(a, b, g) \mapsto (\widetilde{a}, \widetilde{b}, g)$ renders the first order influence function of the approximate functional $\widetilde{\chi}$ equal to the first order influence function of χ evaluated at the projection. Furthermore, the difference between χ and $\widetilde{\chi}$ is quadratic in the distance between $\widetilde{p}$ and p (see (3.6)). The following theorem summarizes the preceding and verifies these properties in the present concrete situation.

THEOREM 6.1. *For given measurable functions $\hat{a}, \underline{a}, \hat{b}, \underline{b}\colon \mathcal{Z} \to \mathbb{R}$ with $\underline{a}$ and $\underline{b}$ bounded away from zero and infinity, define a map $(a, b, g) \mapsto (\widetilde{a}, \widetilde{b}, g)$ by letting $(\widetilde{a} - \hat{a})/\underline{a}$ and $(\widetilde{b} - \hat{b})/\underline{b}$ be the orthogonal projections of $(a - \hat{a})/\underline{a}$ and $(b - \hat{b})/\underline{b}$ in $L_2(\underline{abg})$ onto a closed subspace L. Let $\widetilde{p}$ correspond to $(\widetilde{a}, \widetilde{b}, g)$ and define $\widetilde{\chi}(p) = \chi(\widetilde{p}) + P\chi^{(1)}_{\widetilde{p}}$. Then $\widetilde{\chi}$ has influence function*

$$\widetilde{\chi}^{(1)}_p(X) = A\widetilde{a}(Z)\big(Y - \widetilde{b}(Z)\big) + \widetilde{b}(Z) - \chi\big(\widetilde{p}\big). \tag{6.5}$$

Furthermore, for $\underline{g} = \underline{abg}$,

$$\big|\widetilde{\chi}(p) - \chi(p)\big| \le \Big\|(I - \Pi_p)\frac{\hat{a} - a}{\underline{a}}\Big\|_{2,\underline{g}} \Big\|(I - \Pi_p)\frac{\hat{b} - b}{\underline{b}}\Big\|_{2,\underline{g}}.$$

PROOF. The formula for the influence function agrees with the combination of equations (3.8) and (4.3), and can also be verified directly. In view of (3.5) and (4.4),

$$\widetilde{\chi}(p) - \chi(p) = -\int \big(\widetilde{a} - a\big)\big(\widetilde{b} - b\big)\, g\, d\nu.$$

We rewrite the right side as an integral relative to $\underline{g}\, d\nu$, and next apply the Cauchy-Schwarz inequality. Finally we note that $(\widetilde{a} - a)/\underline{a} = (\widetilde{a} - \hat{a})/\underline{a} - (a - \hat{a})/\underline{a} = (I - \Pi_p)\big((\hat{a} - a)/\underline{a}\big)$, and similarly for b. ∎

The approximation error $\widetilde{\chi}(p) - \chi(p)$ can be rendered arbitrarily small by choosing the space L large enough. Of course, we choose L to be appropriate relative to a-priori assumptions on the functions a and b. If these functions are known to belong to Hölder classes, then L can for instance be chosen as the linear span of the first k basis elements of a suitable orthonormal wavelet basis of $L_2(\nu)$.

To compute higher order influence functions of $\widetilde{\chi}$ we recursively determine influence functions of influence functions, according to the algorithm [1]–[3] in Section 3.1, starting with the influence function of $p \mapsto \widetilde{\chi}_p^{(1)}(x_1) + \chi(\widetilde{p})$, for a fixed x_1. We defer the details of this derivation to Section 10.2, and summarize the result in the following theorem.

To simplify notation, define

$$\begin{aligned} \widetilde{Y} &= A\big(Y - \widetilde{b}(Z)\big)\underline{a}(Z), \\ \widetilde{A} &= \big(A\widetilde{a}(Z) - 1\big)\underline{b}(Z), \\ \underline{A} &= A\underline{a}(Z)\underline{b}(Z). \end{aligned} \tag{6.6}$$

These are the generic variables; indexed versions $\widetilde{Y}_i, \widetilde{A}_i, \underline{A}_i, \ldots$ are defined by adding an index to every variable in the equalities. With this notation and with $\underline{a} = \underline{b} = 1$ the second order influence function (5.3) at $p = \widetilde{p}$ can be written as the symmetrization of $-2\widetilde{Y}_1\Pi_p(Z_1, Z_2)\widetilde{A}_2$. This function was derived in an ad-hoc manner as an approximate or partial influence function of χ, but it is also the exact influence function of $\widetilde{\chi}$. The higher order influence functions of $\widetilde{\chi}$ possess an equally attractive form.

THEOREM 6.2. *An mth order influence function $\widetilde{\chi}_p^{(m)}$ evaluated at $(X_1, \ldots, X_m)$ of the functional $\widetilde{\chi}$ defined in Theorem 6.1 is the degenerate (in $L_2(p)$) part of the variable*

$$(-1)^{j-1} j!\, \widetilde{A}_1 \Pi_{1,2} \underline{A}_2 \Pi_{2,3} \underline{A}_3 \Pi_{3,4} \underline{A}_4 \times \cdots \times \underline{A}_{m-1} \Pi_{m-1,m} \widetilde{Y}_m.$$

Here $\Pi_{i,j}$ is the kernel of the orthogonal projection in $L_2(\underline{abg})$ onto L, evaluated at (Z_i, Z_j).

To obtain the degenerate part of the variable in the preceding lemma, we apply the general formula (2.1) together with Lemma 10.2. Assertions (i) and (ii) of the latter lemma show that the variable is already degenerate relative to X_1 and X_m, while assertion (iii) shows that integrating out the variable X_i for $1 < i < m$ simply collapses $\Pi_{i-1,i}\underline{A}_i\Pi_{i,i+1}$ into $\Pi_{i-1,i+1}$. For

instance, with S_m denoting symmetrization of a function of m variables,

$$
\begin{aligned}
\widetilde{\chi}_p^{(2)}(X_1, X_2) &= -2S_2[\widetilde{A}_1 \Pi_{1,2} \widetilde{Y}_2], \\
\widetilde{\chi}_p^{(3)}(X_1, X_2, X_3) &= 6S_3\Big[\widetilde{A}_1 \Pi_{1,2} \underline{A}_2 \Pi_{2,3} \widetilde{Y}_3 - \widetilde{A}_1 \Pi_{1,3} \widetilde{Y}_3\Big], \\
\widetilde{\chi}_p^{(4)}(X_1, X_2, X_3, X_4) &= -24S_4\Big[\widetilde{A}_1 \Pi_{1,2} \underline{A}_2 \Pi_{2,3} \underline{A}_3 \Pi_{3,4} \widetilde{Y}_4 \\
&\qquad - \widetilde{A}_1 \Pi_{1,3} \underline{A}_3 \Pi_{3,4} \widetilde{Y}_4 - \widetilde{A}_1 \Pi_{1,2} \underline{A}_2 \Pi_{2,4} \widetilde{Y}_4 + \widetilde{A}_1 \Pi_{1,4} \widetilde{Y}_4\Big].
\end{aligned} \tag{6.7}
$$

As shown on the left, but not on the right of the equations, these quantities depend on the unknown parameter $p = (a, b, g)$. In the right sides, the variables $\widetilde{Y}_i$ and $\widetilde{A}_i$ depend on p through $\widetilde{b}$ and $\widetilde{a}$, and hence are not observables. Furthermore, the kernels $\Pi_{i,j}$ depend on g as they are orthogonal projections in $L_2(\underline{abg})$.

7. Parametric rate ($(\alpha + \beta)/2 \geq d/4$). In this section we show that the parameter $\chi(p)$ is estimable at $1/\sqrt{n}$-rate provided the average smoothness $(\alpha + \beta)/2$ is at least $d/4$. We achieve this using the estimator

$$
\hat{\chi}_n = \chi(\hat{p}) + \mathbb{U}_n\Big(\widetilde{\chi}_{\hat{p}}^{(1)} + \tfrac{1}{2}\widetilde{\chi}_{\hat{p}}^{(2)} + \cdots + \frac{1}{m!}\widetilde{\chi}_{\hat{p}}^{(m)}\Big), \tag{7.1}
$$

with the influence functions $\widetilde{\chi}_p^{(j)}$ those of the approximate functional $\widetilde{\chi}$ in Section 6: they are given in Theorems 6.1 and 6.2 for $j = 1$, and $j = 2, \ldots, m$, respectively. (Because the map $p \mapsto \widetilde{p}$ maps $\hat{p}$ into itself, the influence function for $j = 1$ in the display is also the first order influence function (6.5) of of χ, when evaluated at $p = \hat{p}$.)

We assume that the projections Π_p and $\Pi_{\hat{p}}$ map $L_s(\underline{abg})$ to $L_s(\underline{abg})$, for every $s \in \big[r/(r-1), r\big]$, with uniformly bounded norms. (For $r = 2$ this entails only $s = 2$; in this case we define $r/(r-2) = \infty$.)

THEOREM 7.1. *The estimator (7.1), with Π_p a kernel of an orthogonal projection in $L_2(\underline{abg})$ satisfying (12.1) with $\sup_x \Pi_p(x, x) \lesssim k$, satisfies, for a constant c that depends on $\|p/\hat{p}\|_\infty$ only, and $r \geq 2$,*

$$
\begin{aligned}
\hat{\mathrm{E}}_p \hat{\chi}_n - \chi(p) &= O\Big(\|\hat{a} - a\|_r \|\hat{b} - b\|_r \|\hat{g} - g\|_{(m-1)r/(r-2)}^{m-1}\Big) \\
&\qquad + O\Big(\Big\|(I - \Pi_p)\frac{\hat{a} - a}{\underline{a}}\Big\|_2 \Big\|(I - \Pi_p)\frac{\hat{b} - b}{\underline{b}}\Big\|_2\Big), \\
\hat{\mathrm{var}}_p \hat{\chi}_n &\leq \sum_{j=1}^{m} \frac{1}{\binom{n}{j}} c^j k^{j-1}.
\end{aligned}
$$

The first term in the bias is of the order $1+1+(m-1)=m+1$, as to be expected for an estimator based on an mth order influence function; the second term is due to estimating $\widetilde{\chi}$ rather than χ; it is independent of m, and the same as in Theorem 5.1 if $\underline{a}=\underline{b}=1$. The bound on the variance can roughly be understood in that each of the degenerate U-statistics $\mathbb{U}_n\widetilde{\chi}_{\hat{p}}^{(j)}$ in (7.1) contributes a term of order k^{j-1}/n^j.

For α-, β- and γ-regular parameters a,b,g on a d-dimensional domain the range space of the projections Π_p can be chosen so that (5.5) holds and such that there exist estimators $\hat{a},\hat{b},\hat{g}$ of a,b,g, with the first two taking values in this range space, with convergence rates $n^{-\alpha/(2\alpha+d)}$, $n^{-\beta/(2\beta+d)}$ and $n^{-\gamma/(2\gamma+d)}$. Then the second term in the bias (with $\underline{a}=\underline{b}=1$) is of order $(1/k)^{\alpha/d+\beta/d}$. If $(\alpha+\beta)/2\geq d/4$ and we choose $k=n$, then this is of order $1/\sqrt{n}$. For $k=n$ the standard deviation of the resulting estimator is also of the order $1/\sqrt{n}$, while the first term in the bias can be made arbitrarily small by choosing a sufficiently large order m. Specifically, the estimator $\hat{\chi}_n$ attains a $\sqrt{n}$-rate of convergence as soon as

$$m-1\geq\Big(\frac{1}{2}-\frac{\alpha}{2\alpha+d}-\frac{\beta}{2\beta+d}\Big)\Big(\frac{2\gamma+d}{\gamma}\Big). \tag{7.2}$$

For any $\gamma>0$ there exists an order m that satisfies this, and hence the parameter is $\sqrt{n}$-estimable as soon as $(\alpha+\beta)/2\geq d/4$.

More ambitiously, we may aim at attaining the parametric rate for every $\gamma>0$, without a-priori knowledge of γ. This can be achieved if $(\alpha+\beta)/2>d/4$ by using orders $m=m_n$ that increase to infinity with the sample size. In this case the estimator can also be shown to be asymptotically efficient in the semiparametric sense.

THEOREM 7.2. *If $(\alpha+\beta)/2>d/4$, then the estimator (7.1), with $m=\log n$ and Π_p a kernel of an orthogonal projection in $L_2(\underline{abg})$ on a $k=n/(\log n)^2$-dimensional space satisfying (5.5) and (12.1) with $\sup_x\Pi_p(x,x)\lesssim k$, based on preliminary estimators $\hat{a},\hat{b},\hat{g}$ that attain rates $(\log n/n)^{-\delta/(2\delta+d)}$ relative to the uniform norm, satisfies*

$$\sqrt{n}\big(\hat{\chi}_n-\chi(p)-\mathbb{P}_n\widetilde{\chi}_p^{(1)}\big)\xrightarrow{P}0.$$

An estimator that is asymptotically linear in the first order efficient influence function, as in the theorem, is asymptotically optimal in terms of the local asymptotic minimax and convolution theorems (see e.g. [40], Chapter 25). The present estimator $\hat{\chi}_n$ actually looses its efficiency by splitting the sample in a part used to construct the preliminary estimators and a part to form $\mathbb{P}_n$. This can be easily remedied by crossing over the two parts

of the split, and taking the average of the two estimators so obtained. By the theorem these are both asymptotically linear in their sample, and hence their average is asymptotically linear in the full sample and asymptotically efficient.

The proofs of the theorems are deferred to Section 10.3.

8. Minimax rate at lower smoothness ($(\alpha+\beta)/2 < d/4$). If the average a-priori smoothness $(\alpha+\beta)/2$ of the functions a and b falls below $d/4$, then the functional χ cannot be estimated any more at the parametric rate ([29]). The estimator (7.1) of Theorem 7.1 can still be used and, with its bias and variance as given in the theorem properly balanced, attains a certain rate of convergence, faster than the current state-of-the-art linear estimators. However, in this section we present an estimator that is always better, and attains the minimax rate of convergence $n^{-(2\alpha+2\beta)/(2\alpha+2\beta+d)}$ provided that the parameter g is sufficiently regular.

This estimator takes the same general form

$$\hat{\chi}_n = \chi(\hat{p}) + \mathbb{U}_n\Big(\widetilde{\chi}^{(1)}_{\hat{p}} + \tfrac{1}{2}\widetilde{\chi}^{(2)}_{\hat{p}} + \cdots + \frac{1}{m!}\chi^{(m)}_{\hat{p}}\Big), \tag{8.1}$$

as the estimator (7.1), but the influence functions $\chi^{(j)}_p$ for $j \geq 3$ will be different. The idea is to "cut out" certain terms from the influence functions in (7.1) in order to decrease the variance, but without increasing the bias. For clarity we first consider the third order estimator, and next extend to the general mth order. To attain the minimax rate the order m must be fixed to a large enough value so that the first term in the bias given in Theorem 7.1 is no larger than $n^{-(2\alpha+2\beta)/(2\alpha+2\beta+d)}$. (Apart from added complexity there is no loss in choosing m *larger* than needed.)

The third order kernel $\widetilde{\chi}^{(3)}_p$ in (6.7) is the symmetrization of the variable

$$6\widetilde{A}_1\Big(\Pi_p(Z_1,Z_2)\underline{A}_2\Pi_p(Z_2,Z_3) - \Pi_p(Z_1,Z_3)\Big)\widetilde{Y}_3.$$

Here Π_p is the kernel of an orthogonal projection in $L_2(\underline{abg})$ onto a k-dimensional linear space, which we may view as the sum of k projections on one-dimensional spaces. The quantity k^2 in the order $O(k^2/n^3)$ of the variance in Theorem 7.1 for $m = 3$ arises as the number of terms in the product $\Pi_p(Z_1,Z_2)\underline{A}_2\Pi_p(Z_2,Z_3)$ of the two k-dimensional projection kernels. It turns out that this order can be reduced without increasing the bias by cutting out "products of projections on higher base elements".

To make this precise, we partition the projection space in blocks, and

decompose the two projections in the influence function over the blocks:

$$\Pi_p = \sum_{r=0}^{R} \Pi_p^{(k_{r-1},k_r]}, \qquad \Pi_p = \sum_{s=0}^{S} \Pi_p^{(l_{s-1},l_s]}. \tag{8.2}$$

Here $\Pi_p^{(m,n]}$ is the projection on the subspace spanned by base elements with index in intervals $(m,n]$, and $1 = k_{-1} < k_0 < k_1 < \cdots < k_R = k$ and $1 = l_{-1} < l_0 < l_1 < \cdots < l_S = k$ are suitable partitions of the set $\{1,\ldots,k\}$. ("Full" partitions in singleton sets would make the construction conceptual simpler, but a small number of blocks will be needed in our proofs.) The product of the two kernels now becomes a double sum, from which we retain only terms with small values of (r,s). The improved third order influence function is, with as before S_3 denoting symmetrization,

$$\chi_p^{(3)}(X_1,X_2,X_3) = 6S_3\Bigg[\sum_{\substack{(r,s):r+s\le D\\ \vee r=0\vee s=0}}\sum \widetilde{A}_1\Big(\Pi_p^{(k_{r-1},k_r]}(Z_1,Z_2)\underline{A}_2\Pi_p^{(l_{s-1},l_s]}(Z_2,Z_3) - \Pi_p^{(k_{r-1}\vee l_{s-1},k_r\wedge l_s]}(Z_1,Z_3)\Big)\widetilde{Y}_3\Bigg]. \tag{8.3}$$

The negative term in the display is the conditional expectation given Z_1, Z_3 of the leading term, and maintains the degeneracy of the kernel. w For the decomposition (8.2) to be valid, the subspaces corresponding to the blocks must be orthogonal in $L_2(\underline{abg})$. We may achieve this by starting with a standard basis $e_1, \varepsilon_2, \ldots$, with good approximation properties for a target model, and next replacing this by an orthonormal basis in $L_2(\underline{abg})$ by the Gram-Schmidt procedure. For a bounded g the approximation properties will be preserved.

The grids are defined by

$$k_{-1} = 1, \qquad k_r \sim n2^{r/\alpha}, \qquad r = 0,\ldots,R, \tag{8.4}$$

$$l_{-1} = 1, \qquad l_s \sim n2^{s/\beta}, \qquad s = 0,\ldots,S, \tag{8.5}$$

where R and S are chosen such that $k_R \sim l_S \sim k$ (note that $k_0 = l_0 = n$). In these definitions the notation $\sim$ means "equal up to a fixed multiple" (needed to allow that k_r and l_s are (dyadic) integers). For ease of notation let $l_s = l_{-1}$ for $s \le -1$, and $l_s = l_S$ for $s \ge S$.

The grids $k_0 < k_1 < \cdots < k_R$ and $l_0 < l_1 < \cdots < l_S$ partition the integers $n, n+1, \ldots, k$ in R and S groups. As $k_r^\alpha l_s^\beta = 2^{r+s}n^{\alpha+\beta}$, for every $r,s \ge 0$, the cut-off $r+s \le D$ in (8.3) is delimited by the "hyperbola" $i^\alpha j^\beta \sim 2^D n^{\alpha+\beta}$ in the space of indices $(i,j) \in \{1,\ldots,k\}^2$ of base elements used in the two

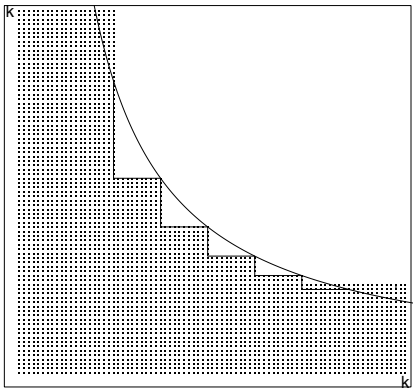

FIG 1. *Both axis carry the indices of the basis functions spanning the projection space L, and point in the plane refers to a product of two projections. Products of projections on pairs of basis functions in the shaded area are included in the third order influence function. The step function refers to the partitions of the indices as in (8.2).*

kernels, with only the pairs below the hyperbola retained (see Figure 1). The intuition behind this hyperbolic cut-off is the product form of the bias (4.4): a higher order correction on the estimator of a may combine with a lower order correction on b, and vice versa, to give an overall correction of the desired order. The overall bias is smaller if the cut-off D is chosen larger, but then more terms are included in the estimator and the variance will be bigger.

Before deriving an optimal value of D, we introduce the mth order estimator for general $m \geq 3$. Again we take the estimator of Theorem 7.1 as starting point, but modify the higher order influence functions $\widetilde{\chi}_p^{(j)}$, for $j = 4, \ldots, m$, similar and in addition to the modification of the third order influence function. For given j the former influence function is given in Theorem 6.2 (with m of the theorem taken equal to j), and is based on a product of $j-1$ projection kernels. We modify this in two steps. For each of the $j-2$ contiguous pairs of kernels $((1st, 2nd), (2nd, 3rd), \ldots, ((j-2)th, (j-1)th))$ we form a new kernel by truncating the pair at the hyperbola as described previously for the third order kernel, and truncating all other kernels at n. Next the modified jth order kernel is the sum of the resulting $j-2$ kernels. More formally, the modified jth order kernel is equal to

$$\chi_p^{(j)}(X_1, \ldots, X_j) = \sum_{i=1}^{j-2} \chi_p^{(j,i)}(X_1, \ldots, X_j), \tag{8.6}$$

where $\chi_p^{(j,i)}(X_1, \ldots, X_j)$ is the symmetrized, degenerate (relative to $L_2(p)$)

part of the variable, for $i = 1, \ldots, j-2$, written in the notation of Theorem 7.1,

$$j!(-1)^{j-1}\,\widetilde{Y}_1\Pi_{1,2}^{(0,n]}\underline{A}_2 \times \cdots \times \underline{A}_{i-1}\Pi_{i-1,i}^{(0,n]}\underline{A}_i \times$$
$$\times \Big[\sum_{\substack{(r,s):r+s\le D\\ \vee r=0\vee s=0}}\sum \Pi_{i,i+1}^{(k_{r-1},k_r]}\underline{A}_{i+1}\Pi_{i+1,i+2}^{(l_{s-1},l_s]}\Big]\underline{A}_{i+2}\Pi_{i+2,i+3}^{(0,n]} \times \cdots \times \underline{A}_{j-1}\Pi_{j-1,j}^{(0,n]}\widetilde{A}_j.$$

For $j = 3$ there is only one pair of kernels, and the construction reduces to the modification (8.3) as discussed previously.

We assume that the projections $\Pi_p^{(0,l]}$ and $\Pi_{\hat{p}}^{(0,l]}$ map $L_s(\underline{abg})$ to $L_s(\underline{abg})$, for every $s \in [r/(r-1), r]$, with uniformly bounded norms.

THEOREM 8.1. *The estimator (8.1) for $m \ge 3$ with the influence functions $\widetilde{\chi}_p^{(j)}$ and $\chi_p^{(j)}$ given in (6.5) and (6.7) for $j = 1, 2$, respectively, and in (8.6) for $j \ge 3$, and with $\Pi_p^{(0,l]}$ kernels of orthogonal projections in $L_2(\underline{abg})$ satisfying (12.1) with $\sup_x \Pi_{\hat{p}}^{(0,l]}(x,x) \lesssim l$, satisfies, for $r \ge 2$ (and $r/(r-2) = \infty$ if $r = 2$),*

$$\hat{\mathrm{E}}_p\hat{\chi}_n - \chi(p) = O\Big(\|\hat{a}-a\|_r\|\hat{b}-b\|_r\|\hat{g}-g\|_{\frac{mr}{r-2}}^{m-1}\Big)$$
$$+ O\Big(\Big\|(I-\Pi_p^{(0,k]})\frac{\hat{a}-a}{\underline{a}}\Big\|_2\Big\|(I-\Pi_p^{(0,k]})\frac{\hat{b}-b}{\underline{b}}\Big\|_2\Big),$$
$$+ O\Big(\sum_{r=1}^{R}\Big\|(I-\Pi_{\hat{p}}^{(0,k_{r-1}]})\Big(\frac{\hat{a}-a}{\underline{a}}\Big)\Big\|_r\Big\|(I-\Pi_{\hat{p}}^{(0,l_{D-r}]})\Big(\frac{\hat{b}-b}{\underline{b}}\Big)\Big\|_r\|\hat{g}-g\|_{\frac{r}{r-2}}\Big)$$
$$+ O\Big(R\Big\|(I-\Pi_{\hat{p}}^{(0,n]})\frac{\hat{a}-a}{\underline{a}}\Big\|_r\Big\|(I-\Pi_{\hat{p}}^{(0,n]})\frac{\hat{b}-b}{\underline{b}}\Big\|_r\|\hat{g}-g\|_{\frac{mr}{r-2}}^2\Big),$$
$$\hat{\mathrm{var}}_p\hat{\chi}_n \lesssim \frac{1}{n} + \frac{k}{n^2} + \frac{D2^{(\frac{1}{\alpha}\vee\frac{1}{\beta})D}}{n}.$$

A proof of the theorem is presented in Sections 10.4 and 10.5.

The first two terms in the bias are the same as in Theorem 7.1; the third and fourth terms are the price paid for cutting out terms from the influence function. The benefit is a reduced variance. We shall show that the boundary parameter D can be chosen such that the third term in the variance (resulting from the third and higher order parts of the influence function) is not bigger than the second term, while the increase in bias is negligible.

Assume that the functions a and b and their estimates are known to belong to models that are well approximated by the base functions $e_1, e_2, \ldots$

in the sense that, for $p \in \{p, \hat{p}\}$, and every value l in one of the two grids (8.4)-(8.5),

$$\Big\|(I - \Pi_p^{(0,l]})\Big(\frac{\hat{a} - a}{\underline{a}}\Big)\Big\|_r \lesssim \Big(\frac{1}{l}\Big)^{\alpha/d}, \tag{8.7}$$

$$\Big\|(I - \Pi_p^{(0,l]})\Big(\frac{\hat{b} - b}{\underline{b}}\Big)\Big\|_r \lesssim \Big(\frac{1}{l}\Big)^{\beta/d}. \tag{8.8}$$

Then the second term in the bias is of the order $(1/k)^{\alpha/d+\beta/d}$, as in Theorem 7.1, which is smaller than the minimax rate $n^{-(2\alpha+2\beta)/(2\alpha+2\beta+d)}$ for

$$k \sim n^{2d/(2\alpha+2\beta+d)}. \tag{8.9}$$

With this choice of k, the upper bound on the variance is of the square minimax rate $n^{-(4\alpha+4\beta)/(2\alpha+2\beta+d)}$ if D is chosen to satisfy

$$2^{(\frac{1}{\alpha} \vee \frac{1}{\beta})D} \sim \frac{1}{\log n} n^{(d-2\alpha-2\beta)/(d+2\alpha+2\beta)}. \tag{8.10}$$

Furthermore, under (8.9) the numbers R, S of grid points are of the order $\log n$.

In the third term of the bias we apply assumptions (8.7)-(8.8) and the identity $k_{r-1}^{\alpha} l_{D-r}^{\beta} \sim n^{\alpha+\beta} 2^D$, which results from (8.4)-(8.5), to see that the third term of the bias is of order

$$\sum_{r=1}^{R} \Big(\frac{1}{k_{r-1}}\Big)^{\alpha/d} \Big(\frac{1}{l_{D-r}}\Big)^{\beta/d} \|\hat{g} - g\|_{r/(r-2)} \le R\Big(\frac{1}{n^{\alpha+\beta} 2^D}\Big)^{1/d} \|\hat{g} - g\|_{r/(r-2)}.$$

If the convergence rate of $\hat{g}$ is $n^{-\gamma/(2\gamma+d)}$, then, for the choice of D given in (8.10), this can (by a calculation) seen to be of smaller order than the minimax rate $n^{-(2\alpha+2\beta)/(2\alpha+2\beta+d)}$ if γ is large enough that

$$\frac{\gamma}{2\gamma + d} > \Big(\frac{\alpha \vee \beta}{d}\Big)\Big(\frac{d - 2\alpha - 2\beta}{d + 2\alpha + 2\beta}\Big). \tag{8.11}$$

The fourth term in the bias can by a similar analysis be seen to be of the order

$$R\Big(\frac{1}{n}\Big)^{\alpha/d} \Big(\frac{1}{n}\Big)^{\beta/d} \|\hat{g} - g\|^2_{(m-2)r/(r-2)}.$$

Again this is smaller than the minimax rate if γ satisfies assumption (8.11).

Finally, if the convergence rates of $\hat{a}$ and $\hat{b}$ are $n^{-\alpha/(2\alpha+d)}$ and $n^{-\beta/(2\beta+d)}$, then the first term in the upper bound of the bias is of the order

$$\Big(\frac{1}{n}\Big)^{\alpha/(2\alpha+d)+\beta/(2\beta+d)+(m-1)\gamma/(2\gamma+d)}.$$

We choose m large enough so that this is of smaller order than the preceding terms. In particular, we can choose it so that this is smaller than the minimax rate.

We summarize this in the following corollary, which is the most advanced result of the paper.

COROLLARY 8.1. *If (8.7)–(8.11) hold, and* $\Pi_p^{(0,l]}$ *are kernels of orthogonal projections in* $L_2(\underline{abg})$ *satisfying (12.1) with* $\sup_x \Pi_{\hat{p}}^{(0,l]}(x,x) \le l$, *then the* m*th order estimator with the kernels (8.6) for* $j \ge 3$ *and sufficiently large* m *and suitable initial estimators, attains the rate* $n^{-(2\alpha+2\beta)/(2\alpha+2\beta+d)}$ *for estimating* $\chi(p)$.

9. Other examples. In this section we briefly indicate a number of other examples for which our general heuristics have been worked out, leading to well known or novel estimators.

9.1. *Density estimation.* Consider estimating a density $\chi(p) = p(a)$ at the fixed point a based on a random sample from p. A first order influence function of this functional would satisfy, for every smooth path $t \mapsto p_t$ with score function g at $t = 0$,

$$\int \chi_p^{(1)} gp \, d\mu = \frac{d}{dt}_{|t=0} \chi(p_t) = g(a)p(a).$$

In a nonparametric situation every zero-mean function g arises as a score function, and hence $\chi_p^{(1)}$ would have to be a "Dirac function at a". Because this does not exist (except for very special p), in this example already a first order influence function fails to exist.

We may approximate the Dirac function by the function $x \mapsto \Pi(a,x)$ for Π the kernel of an orthogonal projection onto a given (large) subspace L of $L_2(\mu)$. Because $\int \Pi(a,x)g(x)p(x)\,d\mu(x) = g(a)p(a)$ for every function g such that $gp \in L$, the function $x \mapsto \Pi(a,x)$ achieves representation for a large set of scores. The corresponding degenerate version is $x \mapsto \Pi(a,x) - \Pi p(a)$, for $\Pi p = \int \Pi(\cdot,x)p(x)\,d\mu(x)$ the projection of p. The corresponding first order estimator (3.1) is

$$\hat{\chi}_n = \chi(\hat{p}_n) + \mathbb{P}_n\big(\Pi(a,\cdot) - \Pi\hat{p}_n(a)\big) = \mathbb{P}_n\Pi(a,\cdot) + \big((I-\Pi)\hat{p}_n\big)(a).$$

If $\hat{p}_n \in L$, then the second term vanishes and the estimator reduces to $\mathbb{P}_n\Pi(a,\cdot)$. This is the usual projection estimator (cf. [25, 32]): if L is spanned by the orthonormal set $e_1, e_2, \ldots, e_k$, then $\Pi(x_1,x_2) = \sum_{i=1}^k e_i(x_1)e_i(x_2)$ and $\hat{\chi}_n = \sum_{i=1}^k (\mathbb{P}_n e_i)e_i(a)$.

Alternative to viewing $x \mapsto \Pi(a, x)$ as an approximation to the "ideal" influence function, we can derive it as the exact influence function of the approximate functional $\widetilde{\chi}(p) = \chi(\Pi p)$.

9.2. *Quadratic functionals.* Consider estimating the functional $\chi(p) = \int p^2 \, d\mu$ based on a random sample of size n from the density p.

The first order influence function of this functional exists on the full nonparametric model, and can be seen to take the form

$$\chi_p^{(1)}(x) = 2\big(p(x) - \chi(p)\big).$$

By the algorithm [1]–[3] of Section 3.1 a second order influence function can be computed as the degenerate part of an influence function of the functional $p \mapsto \bar{\chi}_p^{(1)}(x_1) = 2p(x_1)$, for fixed x_1. As seen in Section 9.1, point evaluation is not a differentiable functional, but has the kernel Π of an orthogonal projection in $L_2(\mu)$ as an *approximate* influence function. Thus an approximate second order influence function of the present functional, minus its projection onto the degenerate functions, is given by

$$\widetilde{\chi}_p^{(2)}(x_1, x_2) = 2\Pi(x_1, x_2) - 2\Pi p(x_1) - 2\Pi p(x_2) + 2\int (\Pi p)^2 \, d\mu.$$

This may also be derived as an exact influence function of the approximate functional $\widetilde{\chi}(p) = \chi(\Pi p)$.

It can be checked that the estimator (7.1) for $m = 2$, given an initial estimator $\hat{p}_n$ that is contained in the range of Π, reduces to $\hat{\chi}_n = \mathbb{U}_n \Pi$, which is a well known estimator ([17]).

9.3. *Doubly robust models.* The heuristics described in Section 3 ought to be applicable in a wide range of estimation problems, but the detailed treatment of the missing data problem in Sections 4–8 shows that their implementation can be involved. Inspection of the proofs reveals that the particular implementation in the latter sections is based on the structure (4.3) of the first order influence function in the missing data problem. The argument extends to semiparametric models with first order influence function of the form

$$\chi_p^{(1)}(x) = a(z)b(z)S_1(x) + a(z)S_2(x) + b(z)S_3(x) + S_4(x) - \chi(p), \tag{9.1}$$

for known functions $S_i(x)$ of the data (i.e. $S = (S_1, S_2, S_3, S_4)$ is a given statistic). The full parameter may be a quadruplet $p \leftrightarrow (a, b, c, f)$, in which f is the marginal density of an observable covariate Z, and c does not appear in (9.1). Other examples of this structure are described in [27, 37].

10. Proofs.

10.1. *Proof of Theorem 5.1.* Write $\hat{\Pi}$ and Π for $\Pi_{\hat{p}}$ and Π_p, respectively, for both the kernels and the corresponding projection operators, and drop p also in $\hat{\mathrm{E}}_p$ and $\hat{\mathrm{var}}_p$. From (4.4) and (5.3) we have

$$\begin{aligned}
&\hat{\mathrm{E}}\hat{\chi}_n - \chi(p) \\
&\quad = -\int (\hat{a}-a)(\hat{b}-b)\, g\, d\nu - \hat{\mathrm{E}} A_1 \big(Y_1 - \hat{b}(Z_1)\big)\big(A_2 \hat{a}(Z_2) - 1\big)\hat{\Pi}(Z_1, Z_2) \\
&\quad = -\int (\hat{a}-a)(\hat{b}-b)\, g\, d\nu + \iint \big[(\hat{a}-a)\times(\hat{b}-b)\big]\,(g\times g)\,\hat{\Pi}\, d\nu \times \nu.
\end{aligned}$$

The double integral on the far right with $\hat{\Pi}$ replaced by Π can be written as the single integral $\int (\hat{a}-a)\Pi(\hat{b}-b)\, g\, d\nu$, for $\Pi(\hat{b}-b)$ the image of $\hat{b}-b$ under the projection Π. Added to the first integral on the right this gives $-\int (\hat{a}-a)(I-\Pi)(\hat{b}-b)\, g\, d\nu$, which is bounded in absolute value by the second term in the upper bound for the bias.

Replacement of $\hat{\Pi}$ by Π in the double integral gives a difference

$$\begin{aligned}
&\iint \big[(\hat{a}-a)\times(\hat{b}-b)\big]\, g\times g\,(\hat{\Pi}-\Pi)\, d\nu\times\nu \\
&\qquad = \int (\hat{a}-a)\left(\hat{\Pi}\Big((\hat{b}-b)\frac{g}{\hat{g}}\Big) - \Pi(\hat{b}-b)\right) g\, d\nu \\
&\qquad \le \|\hat{a}-a\|_s \left\|\hat{\Pi}\Big((\hat{b}-b)\frac{g}{\hat{g}}\Big) - \Pi(\hat{b}-b)\right\|_{r,\hat{g}} \|g/\hat{g}\|_\infty^{1/r},
\end{aligned}$$

by Hölder's inequality, for a conjugate pair (r, s). Considering $\hat{\Pi}$ as the projection in $L_2(\hat{g})$ with weight 1, and Π as the weighted projection in $L_2(\hat{g})$ with weight function $\hat{w} = g/\hat{g}$, we can apply Lemma 12.7(i) (with $q = s/r$ and $rp = s/(s-2)$) to see that this is bounded in absolute value by

$$\|\hat{a}-a\|_s \|\hat{\Pi}\|_{s,\hat{g}} \|\Pi\|_{s,\hat{g}} \|\hat{b}-b\|_{s,\hat{g}} \|\hat{w}-1\|_{s/(s-2),\hat{g}} \|w\|_\infty^{1/r}.$$

Because $\hat{w}$ is assumed bounded away from 0 and infinity, this is of the same order as the first term in the upper bound on the bias (if r replaces s).

Because the function $\chi_{\hat{p}}^{(1)}$ is uniformly bounded, the (conditional) variance of $\mathbb{U}_n \chi_{\hat{p}}^{(1)}$ is of the order $O(1/n)$. Thus for the variance bound it suffices to consider the (conditional) variance of $\mathbb{U}_n \chi_{\hat{p}}^{(2)}$. In view of Lemma 13.1 and (13.1) this is bounded above by a multiple of

$$\left(1 + \left\|\frac{p}{\hat{p}}\right\|_\infty\right)^2 \hat{P}^n (\mathbb{U}_n \chi_{\hat{p}}^{(2)})^2 = \left(1 + \left\|\frac{p}{\hat{p}}\right\|_\infty\right)^2 \binom{n}{2}^{-2} \hat{P}^2 (\chi_{\hat{p}}^{(2)})^2.$$

The variables $A\big(Y-\hat{b}(Z)\big)$ and $\big(A\hat{a}(Z)-1\big)$ are uniformly bounded. Hence the last term on the right is bounded above by a multiple of $n^{-2}\int\hat{\Pi}^2\,(\hat{g}\times\hat{g})\,d\nu\times\nu$, which is equal to k/n^2, by Lemma 12.3.

10.2. *Proof of Theorem 6.2.* We compute the higher order influence functionals of the approximate functional $\widetilde{\chi}$ using algorithm [1]–[3] in Section 3. This starts by computing the second order influence function as the derivative of $p\mapsto\widetilde{\chi}_p^{(1)}(x_1)+\chi(\widetilde{p})$, for fixed x_1. Because the latter functional (given in (6.5)) depends on the parameters only through $\widetilde{a}(z_1)$ and $\widetilde{b}(z_1)$, the following lemma does the main part of the work.

LEMMA 10.1. *For fixed z_1 influence functions of $p\mapsto\widetilde{a}(z_1)$ and $p\mapsto\widetilde{b}(z_1)$ are given by*

$$\begin{aligned}x_2&\mapsto-\underline{a}(z_1)\Pi_p(z_1,z_2)\big(a_2\widetilde{a}(z_2)-1\big)\underline{b}(z_2),\\ x_2&\mapsto\ \underline{b}(z_1)\Pi_p(z_1,z_2)a_2\big(y_2-\widetilde{b}(z_2)\big)\underline{a}(z_2),\end{aligned}$$

where Π_p is the kernel of the orthogonal projection in $L_2(\underline{a}bg)$ onto L.

PROOF. We can write the equation (6.3) determining $\widetilde{a}$ as $\mathrm{E}\big(A\widetilde{a}(Z)-1\big)\underline{b}(Z)l(Z)=0$, for every $l\in L$. Insert a sufficiently regular path p_t, given by parameters (a_t,b_t,f_t), and differentiate the equality relative to t at $t=0$ to find, with γ a score function of the path

$$\mathrm{E}\frac{d}{dt}_{|t=0}\widetilde{a}_t(Z)A\underline{b}(Z)l(Z)=-\mathrm{E}\big(A\widetilde{a}(Z)-1\big)\underline{b}(Z)l(Z)\,\gamma(X).$$

Using the fact that $\mathrm{E}(A|\,Z)=1/a(Z)$, where a is bounded away from zero, we can also write this as

$$\mathrm{E}\frac{\frac{d}{dt}_{|t=0}\widetilde{a}_t(Z)}{\underline{a}(Z)}\frac{(\underline{ab})}{a}(Z)l(Z)=-\mathrm{E}\frac{\big(A\widetilde{a}(Z)-1\big)a(Z)\gamma(X)}{\underline{a}(Z)}\frac{(\underline{ab})(Z)}{a(Z)}l(Z).$$

Because the function $(\widetilde{a}_t-\hat{a})/\underline{a}$ is contained in L for every t by construction, the function $(d/dt)_{|t=0}\widetilde{a}_t/\underline{a}$ is also contained in L. Combined with the validity of the preceding display for every $l\in L$, we conclude that $(d/dt)_{|t=0}\widetilde{a}_t(Z)/\underline{a}(Z)$ is the weighted projection of $-\big(A\widetilde{a}(Z)-1\big)a(Z)\gamma(X)/\underline{a}(Z)$ in $L_2(\mathrm{P})$ onto the space $\{l(Z)\colon l\in L\}$ relative to the weight $(\underline{ab}/a)(Z)$. The projection can be represented in terms of a kernel operator (cf. Lemma 12.1). If $\Pi_p(z_1,z_2)(\underline{ab})(z_2)/a(z_2)$ denotes the kernel, then

$$\begin{aligned}\frac{\frac{d}{dt}_{|t=0}\widetilde{a}_t(z_1)}{\underline{a}(z_1)}&=-\mathrm{E}\Pi_p(z_1,Z_2)\frac{\big(A_2\widetilde{a}(Z_2)-1\big)a(Z_2)\gamma(X_2)}{\underline{a}(Z_2)}\Big(\frac{\underline{ab}}{a}\Big)(Z_2)\\&=-\mathrm{E}\Pi_p(z_1,Z_2)\big(A_2\widetilde{a}(Z_2)-1\big)\underline{b}(Z_2)\,\gamma(X_2).\end{aligned}$$

This represents the derivative on the left as an inner product of the score function γ with the function on the right of the first equation of the lemma (evaluated at X_2). Thus the first assertion of the lemma is proved.

The second assertion is proved similarly. Using that $\mathrm{E}(Y|\,Z) = b(Z)$ and $\mathrm{E}(A|\,Z) = 1/a(Z)$, we start by writing the equation (6.4) defining $\widetilde{b}$ as $\mathrm{E}\big(\widetilde{b}(Z) - Y\big)/\underline{b}(Z)\,(\underline{ab}/a)(Z)l(Z) = 0$, for every $l \in L$. By the same arguments as before we conclude that $(d/dt)_{|t=0}\widetilde{b}_t(Z)$ is the weighted projection of $\big(Y - \widetilde{b}(Z)\big)\gamma(X)/\underline{b}(Z)$ in $L_2(\mathrm{P})$ onto the space $\{l(Z)\colon l \in L\}$, relative to the weight $(\underline{ab}/a)(Z)$. ∎

The first order influence function (6.5) depends on p only through $\widetilde{a}$ and $\widetilde{b}$ and hence the chain rule and the preceding lemma imply that a second order influence function of $\widetilde{\chi}$ is given by the degenerate part of

$$\begin{aligned}\widetilde{\chi}_p^{(2)}(X_1, X_2) = -\Pi_p(Z_1, Z_2)\Big[&A_1\big(Y_1 - \widetilde{b}(Z_1)\big)\underline{a}(Z_1)\big(A_2\widetilde{a}(Z_2) - 1\big)\underline{b}(Z_2)\\ &+ \big(A_1\widetilde{a}(Z_1) - 1\big)\underline{b}(Z_1)A_2\big(Y_2 - \widetilde{b}(Z_2)\big)\underline{a}(Z_2)\Big].\end{aligned} \tag{10.1}$$

(Note that this function is symmetric in (X_1, X_2); Π_p is symmetric, because it is an orthogonal projection kernel.) Actually, this function is already degenerate and hence is the second order influence function of $\widetilde{\chi}$.

LEMMA 10.2. *For any fixed z_1 and z_3,*

(i) $\mathrm{E}_p\Pi_p(z_1, Z_2)\big(A_2\widetilde{a}(Z_2) - 1\big)\underline{b}(Z_2) = 0.$
(ii) $\mathrm{E}_p\Pi_p(z_1, Z_2)A_2\big(Y_2 - \widetilde{b}(Z_2)\big)\underline{a}(Z_2) = 0.$
(iii) $\mathrm{E}_p\Pi_p(z_1, Z_2)A_2(\underline{ab})(Z_2)\Pi_p(Z_2, z_3) = \Pi_p(z_1, z_3).$

PROOF. Because $(\widetilde{a} - \hat{a})(Z)/\underline{a}(Z)$ and $(\widetilde{b} - \hat{b})(Z)/\underline{a}(Z)$ are the weighted projections in $L_2(\mathrm{P})$ of $(a - \hat{a})(Z)/\underline{a}(Z)$ and $\big(Y - \hat{b}(Z)\big)/\underline{b}(Z)$, respectively, onto $\{l(Z)\colon l \in L\}$ relative to the weights $(\underline{ab}/a)(Z)$,

$$\mathrm{E}_{X_2}\Pi_p(Z_1, Z_2)\Big[\frac{\widetilde{a}(Z_2) - \hat{a}(Z_2)}{\underline{a}(Z_2)} - \frac{a(Z_2) - \hat{a}(Z_2)}{\underline{a}(Z_2)}\Big]\frac{\underline{ab}}{a}(Z_2) = 0, \tag{10.2}$$

$$\mathrm{E}_{X_2}\Pi_p(Z_1, Z_2)\Big[\frac{\widetilde{b}(Z_2) - \hat{b}(Z_2)}{\underline{b}(Z_2)} - \frac{Y_2 - \hat{b}(Z_2)}{\underline{b}(Z_2)}\Big](\underline{ab})(Z_2)A = 0. \tag{10.3}$$

These two assertions imply (i) and (ii). The third assertion follows from the fact that Π_p is the kernel of the weighted projection in $L_2(\mathrm{P})$ onto L relative to the weight $(\underline{ab}/a)(Z)$. ∎

The second order influence function (10.1) depends on p through $\widetilde{a}$ and $\widetilde{b}$ and through the kernel Π_p. We proceed to higher orders by differentiating the influence function relative to these components, and applying the chain rule, where we use the influence functions of $p \mapsto \widetilde{a}(x)$ and $p \mapsto \widetilde{b}(x)$ as given previously in Lemma 10.1, and the influence function of $p \mapsto \Pi_p(z_1, z_2)$ as given in Lemma 12.8.

PROOF OF THEOREM 6.2. Denote the symmetrization of the variable in the theorem by $\bar{\chi}_p^{(m)}(X_1, \ldots, X_m)$. Then $\bar{\chi}_p^{(2)}$ is the function $\widetilde{\chi}_p^{(2)}$ given by (10.1), which was seen to be a second order influence function in the preceding discussion. We show by induction on m that $x_{m+1} \mapsto \bar{\chi}_p^{(m+1)}(x_1, \ldots, x_m, x_{m+1})$ is an influence function of $p \mapsto \bar{\chi}_p^{(m)}(x_1, \ldots, x_m)$. The theorem is then a corollary of Lemma 11.2.

By Lemmas 10.1 and 12.8,

(i) The influence function of $p \mapsto \widetilde{Y}_1$ is $x_{m+1} \mapsto -\Pi_p(Z_1, z_{m+1})\underline{A}_1\widetilde{y}_{m+1}$
(ii) The influence function of $p \mapsto \widetilde{A}_1$ is $x_{m+1} \mapsto -\Pi_p(Z_1, z_{m+1})\underline{A}_1\widetilde{a}_{m+1}$.
(iii) The influence function of $p \mapsto \underline{A}_1$ is zero.
(iv) The influence function of $p \mapsto \Pi_p(Z_1, Z_2)$ is $x_{m+1} \mapsto -\Pi_p(Z_1, z_{m+1})\underline{A}_{m+1}\Pi_p(z_{m+1}, Z_2)$.

Applying this repeatedly readily gives an expression for the influence function of $p \mapsto \widetilde{A}_1\Pi_{1,2}\underline{A}_2\Pi_{2,3}\underline{A}_3\Pi_{3,4}\underline{A}_4 \times \cdots \times \underline{A}_{m-1}\Pi_{m-1,m}\widetilde{Y}_m$. The symmetrization of this expression is the same expression, but then with m replaced by $m+1$ and an added minus sign. ∎

10.3. *Proof of Theorems 7.1 and 7.2.* Let $\hat{A}$ and $\hat{Y}$ be $\widetilde{A}$ and $\widetilde{Y}$ as in (6.6) with a and b in their definitions replaced by $\hat{a}$ and $\hat{b}$. Because $\hat{a}$ and $\hat{b}$ are projected onto themselves under the map $(a, b, g) \mapsto (\widetilde{a}, \widetilde{b}, g)$ (see Theorem 6.1), we actually obtain the same variables by replacing $\widetilde{a}$ and $\widetilde{b}$ by $\hat{a}$ and $\hat{b}$, respectively: $\hat{A} = \big(A\hat{a}(Z) - 1\big)\underline{b}(Z)$ and $\hat{Y} = A\big(Y - \hat{b}(Z)\big)\underline{a}(Z)$. Furthermore, let Π and $\hat{\Pi}$ denote the operators Π_p and $\Pi_{\hat{p}}$, respectively, and $\Pi_{i,j}$ and $\hat{\Pi}_{i,j}$ their kernels evaluated at (Z_i, Z_j).

By explicit calculations,

$$\chi(\hat{p}) + \hat{\mathrm{E}}_p\widetilde{\chi}_{\hat{p}}^{(1)} - \chi(p) = -\int (\hat{a} - a)(\hat{b} - b)\, g\, d\nu = \hat{\mathrm{E}}\hat{A}_1\Pi_{1,2}\hat{Y}_2 - \hat{R}, \tag{10.4}$$

for $\hat{R}$ defined by

$$\hat{R} = \int \Big(\frac{\hat{a} - a}{\underline{a}}\Big)(I - \Pi)\Big(\frac{\hat{b} - b}{\underline{b}}\Big)\underline{a}\underline{b}g\, d\nu.$$

The variable $\hat{R}$ is bounded by the second term in the expression for $\hat{\mathrm{E}}_p\hat{\chi}_n - \chi(p)$ in the statement of the theorem. We next show by induction on m that

(10.5)
$$\begin{aligned}&\hat{R} + \chi(\hat{p}) + \hat{\mathrm{E}}\widetilde{\chi}^{(1)}_{\hat{p}} + \cdots + \frac{1}{m!}\hat{\mathrm{E}}\widetilde{\chi}^{(m)}_{\hat{p}} - \chi(p)\\ &= (-1)^{m-1}\hat{\mathrm{E}}\hat{A}_1(\hat{\Pi} - \Pi)_{1,2}\underline{A}_2(\hat{\Pi} - \Pi)_{2,3} \times \cdots \times \underline{A}_{m-1}(\hat{\Pi} - \Pi)_{m-1,m}\hat{Y}_m.\end{aligned}$$

The analysis of the bias can then be concluded by showing that the right side of (10.5) is of the order as the first term given in the theorem.

Equation (10.4) and the definition of $\widetilde{\chi}^{(2)}_p$ readily show that identity (10.5) is true for $m = 2$. We proceed to general m by induction. Relative to its value for m the left side receives for $(m+1)$ the extra term $\hat{\mathrm{E}}\chi^{(m+1)}_{\hat{p}}/(m+1)!$, which is equal to $(-1)^m$ times $\hat{\mathrm{E}}\hat{A}_1\hat{\Pi}_{1,2}\underline{A}_2\hat{\Pi}_{2,3} \times \cdots \times \underline{A}_m\hat{\Pi}_{m,m+1}\hat{Y}_{m+1}$ minus a sum of terms resulting from projections of this leading term. This extra term without the factor $(-1)^m$ (but including the projections) can be written (cf. (6.7) and (2.1))

$$\sum_{i=0}^{m-1}\binom{m-1}{i}\hat{\mathrm{E}}\hat{A}_1\hat{\Pi}_{1,2}\underline{A}_2\hat{\Pi}_{2,3} \times \cdots \times \underline{A}_{m-i}\hat{\Pi}_{m-i,m-i+1}\hat{Y}_{m-i+1}(-1)^i. \tag{10.6}$$

To prove the induction hypothesis for $m + 1$ it suffices to show that this is equal to

$$\begin{aligned}&\hat{\mathrm{E}}\hat{A}_1(\hat{\Pi} - \Pi)_{1,2}\underline{A}_2(\hat{\Pi} - \Pi)_{2,3} \times \cdots \times \underline{A}_{m-1}(\hat{\Pi} - \Pi)_{m-1,m}\hat{Y}_m\\ &\quad + \hat{\mathrm{E}}\hat{A}_1(\hat{\Pi} - \Pi)_{1,2}\underline{A}_2(\hat{\Pi} - \Pi)_{2,3} \times \cdots \times \underline{A}_m(\hat{\Pi} - \Pi)_{m,m+1}\hat{Y}_{m+1}.\end{aligned} \tag{10.7}$$

To achieve this we expand the two terms of the preceding display into sums of expressions of the form, with each $K^{(j)}_{j,j+1}$ equal to $\hat{\Pi}_{j,j+1}$ or $\Pi_{j,j+1}$ and l the number of j for which the first alternative is true,

$$B_l := (-1)^{m-1-l}\,\hat{\mathrm{E}}\hat{A}_1K^{(1)}_{1,2}\underline{A}_2K^{(2)}_{2,3} \times \cdots \times \underline{A}_{m-1}K^{(m-1)}_{m-1,m}\hat{Y}_m, \tag{10.8}$$

and of the same form with $m+1$ replacing m for the second term of (10.7). As the notation suggests the expression in (10.8) depends on l (and m, but this is fixed), but not on which K are equal to $\hat{\Pi}$ or Π. To see this we use that Π is a projection onto L in $L_2(\underline{abg})$, so that $\int \Pi_{1,2}\,\gamma(z_2)\,(\underline{abg})(z_2)\,d\nu(z_2) = \gamma(z_1)$ for every $\gamma \in L$; and $\hat{\Pi}$ is also a projection onto L, so that as a function of one argument $\hat{\Pi}_{1,2}$ is contained in L. This observation yields the identities, for K equal to $\hat{\Pi}$ or Π,

$$\hat{\mathrm{E}}_{Z_j}\Pi_{j-1,j}\underline{A}_jK_{j,j+1} = K_{j-1,j+1} = \hat{\mathrm{E}}_{Z_j}K_{j-1,j}\underline{A}_j\Pi_{j,j+1}.$$

This allows to reduce (10.8) to

$$B_l = (-1)^{m-1-l}\,\hat{\mathrm{E}}\hat{A}_1\hat{\Pi}_{1,2}\underline{A}_2\hat{\Pi}_{2,3} \times \cdots \times \underline{A}_l\hat{\Pi}_{l,l+1}\hat{Y}_{l+1}, \qquad l \geq 1,$$
$$B_0 = (-1)^{m-1}\hat{\mathrm{E}}\hat{A}_1\Pi_{1,2}\hat{Y}_2.$$

Thus after expanding the two terms of (10.7) in the quantities B_l, and simplifying these quantities, we can write their sum (10.7)

$$(B_0 - B_0) + \sum_{l=1}^{m-1}\left(\binom{m}{l} - \binom{m-1}{l}\right) B_l(-1)^{m-l} + B_m.$$

The difference of the binomial coefficients is $\binom{m-1}{l-1}$. The expression is equal to (10.6), as claimed. This completes the proof of (10.5).

Next we bound the right side of (10.5), by taking the expectation in turn with respect to $X_m, X_{m-1}, \ldots, X_1$. For $M_{\hat{w}}$ multiplication by the function $\hat{w} = g/\hat{g}$,

$$\hat{\mathrm{E}}_{X_m}(\hat{\Pi} - \Pi)_{m-1,m}\hat{Y}_m = (\hat{\Pi}M_{\hat{w}} - \Pi)\Big(\frac{\hat{b} - b}{\underline{b}}\Big)(Z_{m-1}).$$

Next, for any function h and $i = m-1, m-2, \ldots, 2$,

$$\hat{\mathrm{E}}_{X_i}(\hat{\Pi} - \Pi)_{i-1,i}\underline{A}_i h(Z_i) = (\hat{\Pi}M_{\hat{w}} - \Pi)h(Z_{i-1}).$$

Combining these equations, we can write the right side of (10.5) in the form

$$(-1)^{m-1}\int \Big(\frac{a - \hat{a}}{\underline{a}}\Big)\Big[\big(\hat{\Pi}M_{\hat{w}} - \Pi\big)^{m-1}\Big(\frac{\hat{b} - b}{\underline{b}}\Big)\Big]\,\underline{abg}\,d\nu.$$

We bound this by first applying Hölder's inequality, with conjugate pair (τ, t) with τ equal to r as in the statement of the theorem, and next Lemma 12.7(iii), with $\hat{\Pi}$ and Π viewed as weighted orthogonal projections in $L_2(\underline{ab}\hat{g})$ with weights 1 and $\hat{w}$, respectively, and $r = \tau(m-1)/(m+\tau-3)$, $p = (m+\tau-3)/(\tau-2)$ and $q = (m+\tau-3)/(m-1)$, so that $rp = (m-1)\tau/(\tau-2)$ and $rq = \tau$ (and m of the lemma taken equal to the present m minus 1).

To bound the (conditional) variance of $\hat{\chi}_n$ we use Lemma 13.1 to see that

$$P^n(\mathbb{U}_n\tilde{\chi}_{\hat{p}}^{(j)})^2 \leq 2j\Big[1 + \Big\|\frac{p}{\hat{p}}\Big\|_\infty\Big]^{2j}\hat{P}^n(\mathbb{U}_n\tilde{\chi}_{\hat{p}}^{(j)})^2 \lesssim \Big[1 + \Big\|\frac{p}{\hat{p}}\Big\|_\infty\Big]^{2j}\frac{2j}{\binom{n}{j}}\hat{P}^j(\tilde{\chi}_{\hat{p}}^{(j)})^2,$$

because $\tilde{\chi}_{\hat{p}}^{(j)}$ is degenerate under $\hat{P}$. The variable $\tilde{\chi}_{\hat{p}}^{(j)}(X_1, \ldots, X_j)/j!$ is the symmetrization of the projection of $\hat{A}_1\hat{\Pi}_{1,2}\underline{A}_2 \cdots \hat{\Pi}_{j-1,j}\hat{Y}_j$ onto the degenerate variables. Because the second moment of a mean of (arbitrary) random

variables is bounded above by the maximum of the second moments of the terms, we can ignore the symmetrization, while the projection decreases the second moment. This shows that

$$\frac{1}{(j!)^2}\hat{P}^j(\tilde{\chi}_{\hat{p}}^{(j)})^2 \le \hat{P}^j(\hat{A}_1\hat{\Pi}_{1,2}\underline{A}_2\cdots\hat{\Pi}_{j-1,j}\hat{Y}_j)^2 \lesssim k^{j-1},$$

by Lemma 12.4 and the assumption that the kernels are bounded by k on the diagonal.

We complete the proof of Theorem 7.1 by bounding the square of $\hat{\chi}_n - \chi(\hat{p})$ by $\sum_{j=1}^m 2^j(\mathbb{U}_n\tilde{\chi}_{\hat{p}}^{(j)}/j!)^2\sum_j 2^{-j}$. The extra factor 2^j can be incorporated in the constant c in the theorem.

For the proof of Theorem 7.2 it clearly suffices to show that

$$\hat{\mathrm{E}}_p\sqrt{n}\big(\hat{\chi}_n - \chi(p) - \mathbb{P}_n\tilde{\chi}_p^{(1)}\big) \xrightarrow{P} 0,$$
$$\hat{\mathrm{var}}_p\sqrt{n}\big(\hat{\chi}_n - \chi(p) - \mathbb{P}_n\tilde{\chi}_p^{(1)}\big) \xrightarrow{P} 0.$$

Because an influence function is centered at mean zero, the first is simply $\sqrt{n}$ times the bias of $\hat{\chi}_n$. By Theorem 7.1 the bias is of the order

$$\Big(\frac{\log n}{n}\Big)^{\alpha/(2\alpha+d)+\beta/(2\beta+d)+\gamma(m-1)/(2\gamma+d)} + \Big(\frac{1}{k}\Big)^{(\alpha+\beta)/d}.$$

The first term is trivially $o(n^{-1/2})$, as $m_n \to \infty$. In the second we write $(\alpha+\beta)/d = r/2$, where $r > 1$ by assumption, and see that it is $o(n^{-1/2})$, since $kn^{-1/r} \to \infty$.

To handle the variance we split the estimator $\hat{\chi}_n$ in its linear and higher order terms. The sum of the variances of the U-statistics of orders 2 to m in $\hat{\chi}_n$ is bounded by the sum of the terms $j \ge 2$ in Theorem 7.1, i.e.

$$\sum_{j=2}^m \frac{c^j k^{j-1}}{\binom{n}{j}} \le \frac{1}{n}\sum_{j=2}^m \Big(\frac{2ckj}{n}\Big)^{j-1}\frac{c2j\sqrt{j}}{e^j},$$

by the inequalities $\binom{n}{j} \ge (n/2)^j/j!$, for $j < n/2$, and $j! \lesssim (j/e)^j\sqrt{j}$, by Stirling's approximation with bound. The expression in brackets is bounded by $2ckm/n \lesssim 1/\log n$, for $m \sim \log n$ and $k \sim n/(\log n)^2$. Thus the sum tends to zero by dominated convergence. Finally the linear term in $\hat{\chi}_n$ gives the contribution

$$\hat{\mathrm{var}}_p\sqrt{n}\big(\mathbb{P}_n\tilde{\chi}_{\hat{p}}^{(1)} - \chi(p) - \mathbb{P}_n\tilde{\chi}_p^{(1)}\big) = \hat{\mathrm{var}}(\tilde{\chi}_{\hat{p}}^{(1)} - \tilde{\chi}_p^{(1)}).$$

From the explicit expression (4.3) for the first order influence function (or (6.5) in the case of $\hat{p}$, which gives an identical function), this is seen to tend to zero by the dominated convergence theorem.

10.4. *Proof of Theorem 8.1 for $m = 3$.* The theorem asserts that the bias of the estimator $\hat{\chi}_n$ is equal to the sum of four terms, the first two of which also arise in the bias of the estimator considered in Theorem 7.1. Therefore, we can prove the assertion on the bias by showing that the expected values of the current estimator $\hat{\chi}_n$ (for $m = 3$) and the estimator in Theorem 7.1 differ by less than the additional bias terms in Theorem 8.1.

The two estimators differ only in their third order influence functions, where the present estimator retains only the terms in the double sum (8.3) with $r = 0$, $s = 0$, or $r + s \leq D$. Thus the difference of the expectations of the two estimators is equal to

$$\sum_{\substack{r+s>D\\ r,s\geq 1}}\sum \hat{\mathrm{E}}_p \hat{A}_1 \Big[\hat{\Pi}^{(k_{r-1},k_r]}(Z_1,Z_2)\underline{A}_2\hat{\Pi}^{(l_{s-1},l_s]}(Z_2,Z_3) - \hat{\Pi}^{(k_{r-1}\vee l_{s-1},k_r\wedge l_s]}(Z_1,Z_3)\Big]\hat{Y}_3.$$

The expectation $\hat{\mathrm{E}}_p$ refers to the variable (X_1, X_2, X_3) for fixed values of the preliminary samples, which are indicated in the "hat" symbols on $\hat{A}_1, \hat{Y}_3$ and the kernels, and hence is an integral relative to the density $(x_1, x_2, x_3) \mapsto p(x_1)p(x_2)p(x_3)$. If we replace $p(x_2)$ in this density by $\hat{p}(x_2)$, then the integral will be zero, as the kernel is degenerate under $\hat{P}$. Thus we may integrate against $(x_1, x_2, x_3) \mapsto p(x_1)(p-\hat{p})(x_2)p(x_3)$. In that case the projection term $\hat{A}_1\hat{\Pi}^{(k_{r-1}\vee l_{s-1},k_r\wedge l_s]}(Z_1, Z_3)\hat{Y}_3$ integrates to zero, as it does not depend on X_2 and $\int (p - \hat{p})(x_2)\, d\mu(x_2) = 0$, and hence can be dropped. Next we condition $\hat{A}_1$ and $\hat{Y}_3$ on Z_1, Z_2, Z_3 and write the preceding display in the form

$$\sum_{\substack{r+s>D\\ r,s\geq 1}}\sum \iiint \frac{\hat{a}-a}{\underline{a}}(z_1)\hat{\Pi}^{(k_{r-1},k_r]}(z_1,z_2)\hat{\Pi}^{(l_{s-1},l_s]}(z_2,z_3)\frac{\hat{b}-b}{\underline{b}}(z_3) \times d\rho(z_1)\, d(\rho-\hat{\rho})(z_2)\, d\rho(z_3).$$

for ρ and $\hat{\rho}$ the measures defined by $d\rho = \underline{abg}\, d\nu$ and $d\hat{\rho} = \underline{ab\hat{g}}\, d\nu$. The double sum can be rewritten as the sum over r running from 1 to R and over s from $D - r + 1$ to S, which gives the equivalent representation, with the $\times$ referring to "tensor products" as explained in Section 2,

$$\sum_{r=1}^{R}\int \Big(\frac{\hat{a}-a}{\underline{a}} \times 1 \times \frac{\hat{b}-b}{\underline{b}}\Big)\Big(\hat{\Pi}^{(k_{r-1},k_r]} \times \hat{\Pi}^{(l_{D-r},k]}\Big)\, d\big(\rho \times (\rho-\hat{\rho}) \times \rho\big).$$

We write $\hat{\Pi}^{(k_{r-1},k_r]} = \hat{\Pi}^{(k_{r-1},k]} - \hat{\Pi}^{(k_r,k]}$, and next arrive at the difference of

two expressions of the type, with $k'_r = k_{r-1}$ and $k'_r = k_r$, respectively,

$$\sum_{r=1}^{R} \int \Big(\frac{\hat{a} - a}{\underline{a}} \times 1 \times \frac{\hat{b} - b}{\underline{b}} \Big) \Big(\hat{\Pi}^{(k'_r,k]} \times \hat{\Pi}^{(l_{D-r},k]} \Big) \, d\big(\rho \times (\rho - \hat{\rho}) \times \rho\big).$$

If the measure of integration were $\hat{\rho} \times (\rho - \hat{\rho}) \times \hat{\rho}$ (with $\hat{\rho}$ instead of ρ), then we could perform the integrals on z_1 and z_3 and next apply Hölder's inequality to bound the resulting expression in absolute value by

$$\sum_{r=1}^{R} \Big\| \hat{\Pi}^{(k'_r,k]} \Big(\frac{\hat{a} - a}{\underline{a}} \Big) \Big\|_r \Big\| \hat{\Pi}^{(l_{D-r},k]} \Big(\frac{\hat{b} - b}{\underline{b}} \Big) \Big\|_r \Big\| \frac{g}{\hat{g}} - 1 \Big\|_{r/(r-2)},$$

where the norms are those of $L_2(\underline{ab}\hat{g})$, which are equivalent to those of $L_2(\nu)$, by assumption. We can write $\hat{\Pi}^{(l,k]} = \hat{\Pi}^{(0,k]}(I - \hat{\Pi}^{(0,l]})$ and use the assumed boundedness of $\hat{\Pi}^{(0,l]}$ as an operator on $L_r(\underline{ab}g)$ to bound this by the third term in the bias.

Replacing $\rho \times (\rho - \hat{\rho}) \times \rho$ by $\hat{\rho} \times (\rho - \hat{\rho}) \times \hat{\rho}$ can be achieved by writing the first and last occurrence of ρ as $\rho = \hat{\rho} + (\rho - \hat{\rho})$ and expanding the resulting expression on the $+$ signs into four terms. One of these has the measure $\hat{\rho} \times (\rho - \hat{\rho}) \times \hat{\rho}$. The other three terms have two or three occurrences of $\rho - \hat{\rho}$, and can be bounded by the first term in the bias (with $m = 3$). This is argued precisely under (10.12) below.

Because the first and second order influence functions are equal to those of the estimator considered in Theorem 7.1, the (conditional) variances of $\mathbb{U}_n \widetilde{\chi}^{(j)}_{\hat{p}}$ for $j = 1, 2$ can be seen to be of the orders $O(1/n)$ and $O(k/n^2)$, respectively, by the same proof. By Lemma 13.1 the variance for $j = 3$ is bounded by (see (13.1))

$$\begin{aligned} &6\Big(1 + \Big\| \frac{p}{\hat{p}} \Big\|_\infty \Big)^6 \hat{P}^n (\mathbb{U}_n \chi^{(3)}_{\hat{p}})^2 \\ &\qquad \lesssim \frac{1}{\binom{n}{3}} \hat{P}^3 \Big(\sum_{r=0}^{R} \hat{A}_1 \hat{\Pi}^{(k_{r-1},k_r]}(Z_1, Z_2) \underline{A}_2 \hat{\Pi}^{(0,l'_{D-r}]}(Z_2, Z_3) \hat{Y}_3 \Big)^2, \end{aligned}$$

where $l'_{D-r} = l_{D-r} \vee n$. After bounding out $\hat{A}_1^2$ and $\hat{Y}_3^2$, we write the squared sum as a double sum. From the fact that the projections $\hat{\Pi}^{(k_{r-1},k_r]}$ are orthogonal for different r, it follows that the off-diagonal terms of the double sum vanish (the expectation with respect to X_1 is zero). Thus the preceding display is bounded above by a multiple of

$$\frac{1}{n^3} \sum_{r=0}^{R} \hat{P}^3 \big(\hat{\Pi}^{(k_{r-1},k_r]}(Z_1, Z_2) \underline{A}_2 \hat{\Pi}^{(0,l'_{D-r}]}(Z_2, Z_3) \big)^2.$$

By Lemmas 12.4 and 12.3 and the assumption that $\sup_z \hat{\Pi}^{(0,l]}(z,z) \lesssim l$ this is bounded by a multiple of

$$\frac{1}{n^3}\sum_{r=0}^{R}(k_r - k_{r-1})l'_{D-r} \le \frac{1}{n^3}\Big(nk + \sum_{r=1}^{R}(k_r - k_{r-1})(l_{D-r} + n)\Big).$$

By (8.4) $k_r - k_{r-1} = (1 - 2^{-\alpha})k_r \lesssim k_r = n2^{r/\alpha}$ for $r \ge 1$. On substituting this in the display, and noting that $l_{D-r} = 0$ if $r > D$, we see that this is bounded $k/n^2 + 2^{D/\alpha \vee D/\beta}/n$ if $\alpha \ne \beta$ and bounded by $k/n^2 + D2^{D/\alpha}/n$ if $\alpha = \beta$.

SUPPLEMENTARY MATERIAL

Supplement: ***Higher order estimating equations*** (doi: COMPLETED BY THE TYPESETTER; .pdf). The remainder of the paper is given in the supplement.

REFERENCES

[1] Bickel, P. J. (1982). On adaptive estimation. *Ann. Statist.* **10**, 3, 647–671. MR663424 (84a:62045)

[2] Bickel, P. J., Klaassen, C. A. J., Ritov, Y., and Wellner, J. A. (1993). *Efficient and adaptive estimation for semiparametric models.* Johns Hopkins Series in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD. MR1245941 (94m:62007)

[3] Bickel, P. J. and Ritov, Y. (1988). Estimating integrated squared density derivatives: sharp best order of convergence estimates. *Sankhyā Ser. A* **50**, 3, 381–393. MR1065550 (91e:62079)

[4] Birgé, L. and Massart, P. (1995). Estimation of integral functionals of a density. *Ann. Statist.* **23**, 1, 11–29. MR1331653 (96c:62065)

[5] Bolthausen, E., Perkins, E., and van der Vaart, A. (2002). *Lectures on probability theory and statistics.* Lecture Notes in Mathematics, Vol. **1781**. Springer-Verlag, Berlin. Lectures from the 29th Summer School on Probability Theory held in Saint-Flour, July 8–24, 1999, Edited by Pierre Bernard. MR1915443 (2003d:60004)

[6] Cai, T. T. and Low, M. G. (2005). Nonquadratic estimators of a quadratic functional. *Ann. Statist.* **33**, 6, 2930–2956. http://dx.doi.org/10.1214/009053605000000147. MR2253108 (2007k:62058)

[7] Cai, T. T. and Low, M. G. (2006). Optimal adaptive estimation of a quadratic functional. *Ann. Statist.* **34**, 5, 2298–2325. http://dx.doi.org/10.1214/009053606000000849. MR2291501 (2008m:62054)

[8] Cohen, A., Dahmen, W., Daubechies, I., and DeVore, R. (2001). Tree approximation and optimal encoding. *Appl. Comput. Harmon. Anal.* **11**, 2, 192–226. http://dx.doi.org/10.1006/acha.2001.0336. MR1848303 (2002g:42048)

[9] Daubechies, I. (1992). *Ten lectures on wavelets.* CBMS-NSF Regional Conference Series in Applied Mathematics, Vol. **61**. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA. MR1162107 (93e:42045)

[10] Donoho, D. L. and Nussbaum, M. (1990). Minimax quadratic estimation of a quadratic functional. *J. Complexity* **6**, 3, 290–323. http://dx.doi.org/10.1016/0885-064X(90)90025-9. MR1081043 (91m:65343)

[11] Hampel, F. R., Ronchetti, E. M., Rousseeuw, P. J., and Stahel, W. A. (1986). *Robust statistics*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons, Inc., New York. The approach based on influence functions. MR829458 (87k:62054)

[12] Härdle, W., Kerkyacharian, G., Picard, D., and Tsybakov, A. (1998). *Wavelets, approximation, and statistical applications*. Lecture Notes in Statistics, Vol. **129**. Springer-Verlag, New York. MR1618204 (99f:42065)

[13] Has′minskiĭ, R. Z. and Ibragimov, I. A. (1979). On the nonparametric estimation of functionals. In *Proceedings of the Second Prague Symposium on Asymptotic Statistics (Hradec Králové, 1978)*. North-Holland, Amsterdam-New York, 41–51. MR571174 (81j:62076)

[14] Huber, P. J. and Ronchetti, E. M. (2009). *Robust statistics*, Second ed. Wiley Series in Probability and Statistics. John Wiley & Sons, Inc., Hoboken, NJ. http://dx.doi.org/10.1002/9780470434697. MR2488795 (2010j:62004)

[15] Kerkyacharian, G. and Picard, D. (1996). Estimating nonquadratic functionals of a density using Haar wavelets. *Ann. Statist.* **24**, 2, 485–507. MR1394973 (97e:62062)

[16] Koševnik, J. A. and Levit, B. J. (1976). On a nonparametric analogue of the information matrix. *Teor. Verojatnost. i Primenen.* **21**, 4, 759–774. MR0428578 (55 #1599)

[17] Laurent, B. (1997). Estimation of integral functionals of a density and its derivatives. *Bernoulli* **3**, 2, 181–211. MR1466306 (99c:62144)

[18] Laurent, B. and Massart, P. (2000). Adaptive estimation of a quadratic functional by model selection. *Ann. Statist.* **28**, 5, 1302–1338. MR1805785 (2002c:62052)

[19] Lindsay, B. G. (1983). Efficiency of the conditional score in a mixture setting. *Ann. Statist.* **11**, 2, 486–497. MR696061 (84h:62050)

[20] Murphy, S. A. and van der Vaart, A. W. (2000). On profile likelihood. *J. Amer. Statist. Assoc.* **95**, 450, 449–485. With comments and a rejoinder by the authors. MR1803168 (2002a:62143)

[21] Nemirovski, A. (2000). Topics in non-parametric statistics. In *Lectures on probability theory and statistics (Saint-Flour, 1998)*. Lecture Notes in Math., Vol. **1738**. Springer, Berlin, 85–277. MR1775640 (2001h:62074)

[22] Pfanzagl, J. (1982). *Contributions to a general asymptotic statistical theory*. Lecture Notes in Statistics, Vol. **13**. Springer-Verlag, New York. With the assistance of W. Wefelmeyer. MR675954 (84i:62036)

[23] Pfanzagl, J. (1985). *Asymptotic expansions for general statistical models*. Lecture Notes in Statistics, Vol. **31**. Springer-Verlag, Berlin. With the assistance of W. Wefelmeyer. MR810004 (87i:62004)

[24] Pfanzagl, J. (1990). *Estimation in semiparametric models*. Lecture Notes in Statistics, Vol. **63**. Springer-Verlag, New York. Some recent developments. MR1048589 (91f:62074)

[25] Prakasa Rao, B. L. S. (1983). *Nonparametric functional estimation*. Probability and Mathematical Statistics. Academic Press Inc. [Harcourt Brace Jovanovich Publishers], New York. MR740865 (86m:62076)

[26] Robins, J., Li, L., Tchetgen, E., and van der Vaart, A. Supplement to "higher order estimating equations for high-dimensional models".

[27] Robins, J., Li, L., Tchetgen, E., and van der Vaart, A. (2008). Higher order influence functions and minimax estimation of nonlinear functionals. In *Probability and statistics: essays in honor of David A. Freedman*. Inst. Math. Stat. Collect., Vol. **2**. Inst. Math. Statist., Beachwood, OH, 335–421. http://dx.doi.org/10.1214/193940307000000527. MR2459958 (2010b:62115)

[28] Robins, J., Li, L., Tchetgen, E., and van der Vaart, A. (2009a). Quadratic semiparametric von mises calculus. *Metrika 69*, 227–247.

[29] Robins, J., Li, L., Tchetgen, E., and van der Vaart, A. (2009b). Semiparametric minimax rates. *Electron. J. Stat. 3*, 1305–1321.

[30] Robins, J. M. and Rotnitzky, A. (1995). Semiparametric efficiency in multivariate regression models with missing data. *J. Amer. Statist. Assoc.* **90**, 429, 122–129. MR1325119 (96d:62084)

[31] Rotnitzky, A. and Robins, J. M. (1995). Semi-parametric estimation of models for means and covariances in the presence of missing data. *Scand. J. Statist.* **22**, 3, 323–333. MR1363216 (96j:62090)

[32] Tsybakov, A. B. (2004). *Introduction à l'estimation non-paramétrique*. Mathématiques & Applications (Berlin) [Mathematics & Applications], Vol. **41**. Springer-Verlag, Berlin. MR2013911 (2005a:62007)

[33] v. Mises, R. (1947). On the asymptotic distribution of differentiable statistical functions. *Ann. Math. Statistics 18*, 309–348. MR0022330 (9,194h)

[34] van der Laan, M. J. and Robins, J. M. (2003). *Unified methods for censored longitudinal data and causality*. Springer Series in Statistics. Springer-Verlag, New York. MR1958123 (2003m:62003)

[35] van der Vaart, A. (1991). On differentiable functionals. *Ann. Statist.* **19**, 1, 178–204. MR1091845 (92i:62100)

[36] Van der Vaart, A. (1996). Efficient maximum likelihood estimation in semiparametric mixture models. *Ann. Statist.* **24**, 2, 862–878. http://dx.doi.org/10.1214/aos/1032894470. MR1394993 (97d:62096)

[37] van der Vaart, A. (2014). Higher Order Tangent Spaces and Influence Functions. *Statist. Sci.* **29**, 4, 679–686. http://dx.doi.org/10.1214/14-STS478. MR3300365

[38] van der Vaart, A. W. (1988a). Estimating a real parameter in a class of semiparametric models. *Ann. Statist.* **16**, 4, 1450–1474. http://dx.doi.org/10.1214/aos/1176351048. MR964933 (89m:62032)

[39] van der Vaart, A. W. (1988b). *Statistical estimation in large parameter spaces*. CWI Tract, Vol. **44**. Stichting Mathematisch Centrum Centrum voor Wiskunde en Informatica, Amsterdam. MR927725 (89e:62049)

[40] van der Vaart, A. W. (1998). *Asymptotic statistics*. Cambridge Series in Statistical and Probabilistic Mathematics, Vol. **3**. Cambridge University Press, Cambridge. MR1652247 (2000c:62003)

[41] van der Vaart, A. W. and Wellner, J. A. (1996). *Weak convergence and empirical processes*. Springer Series in Statistics. Springer-Verlag, New York. With applications to statistics. MR1385671 (97g:60035)

[42] van Zwet, W. R. (1984). A Berry-Esseen bound for symmetric statistics. *Z. Wahrsch. Verw. Gebiete* **66**, 3, 425–440. http://dx.doi.org/10.1007/BF00533707. MR751580 (86h:60063)

[43] Waterman, R. P. and Lindsay, B. G. (1996). Projected score methods for approximating conditional scores. *Biometrika* **83**, 1, 1–13. http://dx.doi.org/10.1093/biomet/83.1.1. MR1399151 (98g:62044)

Submitted to the Annals of Statistics
arXiv: math.PR/0000000

SUPPLEMENT TO "HIGHER ORDER ESTIMATING EQUATIONS FOR HIGH-DIMENSIONAL MODELS"

By James Robins, Lingling Li, Rajarshi Mukherjee, Eric Tchetgen Tchetgen and Aad van der Vaart

Harvard University and Universiteit Leiden

This supplement contains the proof of Theorem 8.1 in the case that $m > 3$, and it contains three appendices.

10.5. *Proof of Theorem 8.1.* As in the proof for $m = 3$ it suffices to compare the bias with the bias of the estimator in Theorems 7.1. In the estimator of order $m > 5$ not every of the additional bias terms of orders $j = 4, \ldots, m-1$ is individually small, but the sum is small due to a cancellation among these terms. The analysis therefore requires careful bookkeeping, for which we introduce the following notation.

A string $\delta^1\delta^2\cdots\delta^{j-1}$ of symbols $\delta^i \in \{0, 1, \text{-}\}$ refers to an expectation of a variable

$$\hat{A}_1\hat{\Pi}_{1,2}^{\delta_1}\underline{A}_2\hat{\Pi}_{2,3}^{\delta^2}\underline{A}_3\cdots\underline{A}_{j-1}\hat{\Pi}_{j-1,j}^{\delta^j}\hat{Y}_j, \tag{10.9}$$

where

$$\hat{\Pi}_{i,j}^{\delta} = \begin{cases} \Pi_{\hat{p}}^{(0,n]}(Z_i, Z_j), & \text{if } \delta = 0, \\ \Pi_{\hat{p}}^{(n,k]}(Z_i, Z_j), & \text{if } \delta = 1, \\ \Pi_{\hat{p}}^{(0,k]}(Z_i, Z_j), & \text{if } \delta = \text{-}. \end{cases}$$

Furthermore, a string $\delta^1\cdots\delta^{i-1}\overline{H}\delta^{i+2}\cdots\delta^{j-1}$ refers to the expectation of a variable

$$\hat{A}_1\hat{\Pi}_{1,2}^{\delta_1}\cdots\hat{\Pi}_{i-1,i}^{\delta^{i-1}}\underline{A}_i\Big(\sum_{\substack{r+s>D\\ r,s\geq 1}}\sum\hat{\Pi}_{i,i+1}^{(k_{r-1},k_r]}\underline{A}_{i+1}\hat{\Pi}_{i+1,i+2}^{(l_{s-1},l_s]}\Big)\underline{A}_{i+2}\hat{\Pi}_{i+2,i+3}^{\delta^{i+2}}\cdots\hat{\Pi}_{j-1,j}^{\delta^j}\hat{Y}_j.$$

Thus the symbol $\overline{H}$ refers to terms in pairs of projection kernels above the hyperbola, as involved in the construction of the estimator. Similarly,

[†]The research leading to these results has received funding from the European Research Council underERC GrantAgreement320637.

AMS 2000 subject classifications: Primary 62G05, 62G20, 62G20, 62F25
Keywords and phrases: Nonlinear functional, nonparametric estimation, U-statistic, influence function, tangent space

we let the same strings but with the symbol H instead of $\overline{H}$ refer to the complementary terms below the hyperbola, but above n.

Every of these strings will stand for an expected value; for the first and last variables X_1 and X_j this is computed relative to p, but for the middle variables $X_2, \dots, X_{j-1}$ this is relative to $p - \hat{p}$. We add further notation for expectations on X_1 and X_j taken relative to $p - \hat{p}$ or $\hat{p}$, by preceding (for X_1) or succeeding (for X_j) the string with a subscript $_d$ (for the difference $p - \hat{p}$) or $_\wedge$ (for $\hat{p}$). This gives, for instance,

$$01\text{-}1_\wedge = \int \frac{a - \hat{a}}{\underline{a}}(z_1)\hat{\Pi}_{1,2}^{(0,n]}\hat{\Pi}_{2,3}^{(n,k]}\hat{\Pi}_{3,4}^{(0,k]}\hat{\Pi}_{4,5}^{(n,k]}\frac{b - \hat{b}}{\underline{b}}(z_5) \times d\rho(z_1)\, d\prod_{i=2}^{4}(\rho - \hat{\rho})(z_i)\, d\hat{\rho}(z_5),$$

$${}_d1\text{-}1_\wedge = \int \frac{a - \hat{a}}{\underline{a}}(z_1)\hat{\Pi}_{1,2}^{(n,k]}\hat{\Pi}_{2,3}^{(0,k]}\hat{\Pi}_{3,4}^{(n,k]}\frac{b - \hat{b}}{\underline{b}}(z_4)\, d\prod_{i=1}^{3}(\rho - \hat{\rho})(z_i)\, d\hat{\rho}(z_4).$$

The notations $d\rho = \underline{abg}\, d\nu$, $d\hat{\rho} = \underline{ab}\hat{g}\, d\nu$, and $\hat{\Pi}_{i,j}$ are as in the proof of the theorem for $m = 3$, and the (five- and four-fold) integrals arise after conditioning (10.9) on the variables Z_i, as in the same proof.

The jth order kernel of the estimator in Theorem 7.1 corresponds to the product of $j-1$ projection kernels with ranges $(0, k]$, and is represented by a string of $j-1$ dashes: $\text{-}\cdots\text{-}$. To construct the estimator of Theorem 8.1 we partition the range of a single kernel as $(0, k] = (0, n] \cup (n, k]$, or the range of a contiguous pair of kernels as $(0, k]^2 = (0, n]^2 \cup H \cup \overline{H}$. By expanding the corresponding product of sums of two or three (pairs of) kernels we obtain a decomposition of $\text{-}\cdots\text{-}$ into sequences with symbols $0, 1, \text{-}, H, \overline{H}$. The terms retained in the estimator of Theorem 8.1 are represented by the sequences $\delta^1 \dots \delta^{j-1} \in \{0, 1, \text{-}\}^{j-1}$ with $j-1$ or $j-2$ symbols 0, and the sequences $H0\cdots0, 0H0\cdots0, 0\cdots0H$. All other terms are left out; for instance, for $j = 3, 4, 5, 6, 7$ the terms that are left out are given by the

sequences

$j=3$	$j=4$	$j=5$	$j=6$	$j=7$
$\overline{H}$	1-1	1--1	1---1	1----1
	$\overline{H}0$	01-1	01--1	01---1
	$0\overline{H}$	1-10	1--10	1---10
		$\overline{H}00$	001-1	001--1
		$0\overline{H}0$	01-10	01--10
		$\overline{H}00$	1-100	1--100
			$000\overline{H}$	0001-1
			$00\overline{H}0$	001-10
			$0\overline{H}00$	01-100
			$\overline{H}000$	1-1000
				$0000\overline{H}$
				$00\overline{H}00$
				$0\overline{H}000$
				$\overline{H}0000$

In this table the strings are categorized by the numbers of 0s on their left and right sides. The nonzero middle part of a string always has the form $1\cdots\text{-}\cdots 1$, with at least one -, or $\overline{H}$, which may be considered as taking the place of 11.

We claim that the difference of the biases of the estimators in Theorems 7.1 and 8.1 is the alternating (on the order) sum of these strings (or, rather, of the expectations they represent). For instance, the extra bias for $m=5$ is equal to the sum of all strings in the table under $j=5$ minus the strings under $j=4$ plus the string under $j=3$.

To see this we note first that the leading factorial $j!$ in the definition of the jth order influence function $\widetilde{\chi}_p^{(j)}$ in Theorem 6.2 and its reduced version $\chi_p^{(j)}$ in Theorem 8.1 cancels the factorial in the definition of the estimator (8.1), while the factor $(-1)^{j-1}$ causes alternation of signs between the orders. The extra bias is the sum over j of the expectation under P^j of the sum of the terms left out of the jth influence function. Because by its construction the influence function is degenerate relative to $X_2,\dots,X_{j-1}$ with respect to $\hat{P}$, the expectation can be equivalently taken relative to $p-\hat{p}$ for these variables. Following this substitution, the projection of the leading term (10.9), which creates the degeneracy, can be dropped, and the expectation reduces to a number as represented by one of the strings $\delta^1\dots\delta^{j-1}$ or $\cdots\overline{H}\cdots$ introduced previously. This last reasoning is similar as in the proof for $m=3$, where the projection is shown explicitly.

We proceed to bound the alternating sum of the "left-out strings". There is cancellation of expectations between terms that are one or two orders

apart, i.e. of strings that differ by one or two symbols. The relevant reduction formulas are, for any $\varepsilon, \delta \in \{0, 1, \text{-}, \overline{H}\}$ and any intermediate symbols $\cdots$,

$$
\begin{aligned}
{}_\wedge 0\varepsilon\cdots\delta 0_\wedge &= {}_\wedge 0\varepsilon\cdots\delta_d + {}_d\varepsilon\cdots\delta 0_\wedge - {}_d\varepsilon\cdots\delta_d + {}_\wedge\bar{1}\varepsilon\cdots\delta\bar{1}_\wedge,\\
{}_\wedge\varepsilon\cdots\delta 0_\wedge &= {}_\wedge\varepsilon\cdots\delta_d - {}_\wedge\varepsilon\cdots\delta\bar{1}_\wedge,\\
{}_\wedge 0\varepsilon\cdots\delta_\wedge &= {}_d\varepsilon\cdots\delta_\wedge - {}_\wedge\bar{1}\varepsilon\cdots\delta_\wedge.
\end{aligned}
\tag{10.10}
$$

The rightmost term of each formula is a remainder term, which we may view as being defined by the formula. The idea of the equations is to remove a symbol 0 at the beginning or end of a string with its mark ${}_\wedge$, where the new, shorter string carries mark ${}_d$. (The second and third formulas, even though valid, will be used only with ε or δ unequal to 0, respectively.) The first formula is true if the remainder string ${}_\wedge\bar{1}\varepsilon\cdots\delta\bar{1}_\wedge$ is interpreted as

$$
\begin{aligned}
\int \Big[(\hat{\Pi}^{(0,n]} - I)\frac{a-\hat{a}}{\underline{a}}\Big](z_2)\,\hat{\Pi}^{\varepsilon}_{2,3}\cdots\hat{\Pi}^{\delta}_{j-2,j-1}\times\\
\times\Big[(\hat{\Pi}^{(0,n]} - I)\frac{b-\hat{b}}{\underline{b}}\Big](z_{j-1})\,d\prod_{i=2}^{j-1}(\rho-\hat{\rho})(z_i).
\end{aligned}
$$

Indeed, the four expectations obtained by expanding this last integral on the minus signs in the two appearances of $\hat{\Pi}^{(0,n]} - I$ are the four strings ${}_\wedge 0\varepsilon\cdots\delta 0_\wedge$, ${}_\wedge 0\varepsilon\cdots\delta_d$, ${}_d\varepsilon\cdots\delta 0_\wedge$, and ${}_d\varepsilon\cdots\delta_d$ in the first reduction formula, with positive, negative, negative and positive signs. This follows by integrating the latter strings on the first and/or last variables, and using identities such as $\int \alpha(z_1)\hat{\Pi}^{(0,n]}_{1,2}\,d\hat{\rho}(z_1) = \hat{\Pi}^{(0,n]}\alpha(z_2)$. The second and third formulas are obtained similarly, and more easily, with the appropriate definitions of the remainder strings. (The notation $\bar{1}$ is motivated by the fact that $\int \alpha(z_1)\hat{\Pi}^{(n,k]}_{1,2}\,d\hat{\rho}(z_1)$, which is represented by a 1, is equal to $\hat{\Pi}^{(n,k]}\alpha(z_2)$, which is $(I - \hat{\Pi}^{(0,n]})\alpha(z_2)$ up to terms "above k".)

We now proceed in two steps to rewrite all strings that make up the difference of the influence functions for $j = 4, \ldots, m$. First we write $p(x_1) = \hat{p}(x_1) + (p - \hat{p})(x_1)$ and similarly for the density of X_j, and expand on the plus signs, to rewrite every string $\varepsilon\cdots\delta$ as:

$$
\varepsilon\cdots\delta = {}_\wedge\varepsilon\cdots\delta_\wedge + {}_\wedge\varepsilon\cdots\delta_d + {}_d\varepsilon\cdots\delta_\wedge + {}_d\varepsilon\cdots\delta_d.
\tag{10.11}
$$

Second, if one or both of ε and δ are 0, we expand the first string on the right side (with two ${}_\wedge$) using the reduction formulas (10.10), where we use the first formula if both $\varepsilon = \delta = 0$ and the second or third if one of ε, δ is 0. After doing this for all strings up to some order, we end up with:

(i) strings of the type ${}_\wedge\varepsilon\cdots\delta_\wedge$, with both $\varepsilon, \delta \in \{1, \overline{H}, \bar{1}\}$.

(ii) strings of the type ${}_\wedge\varepsilon\cdots\delta_d$.
(iii) strings of the type ${}_d\varepsilon\cdots\delta_\wedge$.
(iv) strings of the type ${}_d\varepsilon\cdots\delta_d$.

Note that the reduction formulas (10.10), as we applied them, produce a string ${}_\wedge\varepsilon\cdots\delta_\wedge$ with two $\wedge$ only if $\varepsilon\cdots\delta$ is of type (i). We shall show that the strings of type (i) are individually small, while the contributions of the other types are small after cancellation.

Strings of type (i) can be bounded with the help of Lemma 12.6. For $\varepsilon\cdots\delta$ a string of length $j-3$ with $j\geq 4$, with $\overline{H}$ counted as having length 2,

$$\left|{}_\wedge 1\varepsilon\cdots\delta 1_\wedge\right| \leq \left\|\hat{\Pi}^{(n,k]}\frac{a-\hat{a}}{\underline{a}}\right\|_r \left\|\hat{\Pi}^{(n,k]}\frac{b-\hat{b}}{\underline{b}}\right\|_r \|g-\hat{g}\|^{j-2}_{(j-2)r/(r-2)}.$$

For $\bar{1}$ instead of 1 the similar statement is true, but with $I-\hat{\Pi}^{(0,n]}$ replacing $\hat{\Pi}^{(n,k]}$. Since the projections are norm-decreasing up to a constant by assumption (and Lemma 12.2), the norms in the display are bounded up to a constant by the norms of the functions $(I-\hat{\Pi}^{(0,n]})\alpha$, for $\alpha=(a-\hat{a})/\underline{a}$ or $\alpha=(b-\hat{b})/\underline{b}$, respectively. Strings of type (i) starting or ending with $\overline{H}$ and at least one other symbol can be treated in the same manner, as the kernels in $\overline{H}$ all start above n. These strings minimally give a *square* estimation norm $\|g-\hat{g}\|^2_{(m-2)r/(r-2)}$, and are accounted for in the fourth term of the bound on the bias in Theorem 8.1. The only string of type (i) of length 2 is $\overline{H}$. In the proof for $m=3$ this was shown to be accounted for by the third term of the bias bound.

Every string of type (ii) arises both from the initial expansion (10.11) of $\varepsilon\cdots\delta$, and from the secondary expansion (10.10) of ${}_\wedge\varepsilon\cdots\delta 0_\wedge$. In the expansions they carry the same sign, but as they arise at different orders, the alternation of signs in the orders makes them cancel. The same analysis applies to strings of type (iii). Finally, strings of type (iv) arise from the initial expansion of the string $\varepsilon\cdots\delta$ and from the secondary expansion of the string ${}_\wedge 0\varepsilon\cdots\delta 0_\wedge$, with the opposite sign. As the latter strings arise at orders that differ by two, they also cancel.

If we consider terms up to order m, then the strings that cancel versus strings at orders $m+1$ or $m+2$ are left. These are the strings ${}_d\varepsilon\cdots\delta_d$ of length $m-2$ and $m-1$, with $\overline{H}$ counted as two symbols, and the strings ${}_d\varepsilon\cdots\delta_\wedge$ and ${}_\wedge\varepsilon\cdots\delta_d$ of length $m-1$. In view of Lemma 12.6, for $\varepsilon\cdots\delta$ of

length $j-1$,

$$|\,{}_d\varepsilon\cdots\delta_d| \lesssim \left\|\hat{\Pi}^{\varepsilon}\Big(\frac{a-\hat{a}}{\underline{a}}\frac{g-\hat{g}}{\hat{g}}\Big)\right\|_s \left\|\hat{\Pi}^{\delta}\Big(\frac{b-\hat{b}}{\underline{b}}\frac{g-\hat{g}}{\hat{g}}\Big)\right\|_s \|g-\hat{g}\|^{j-2}_{(j-2)s/(s-2)}$$

$$\lesssim \left\|\frac{a-\hat{a}}{\underline{a}}\right\|_{sp}\left\|\frac{g-\hat{g}}{\hat{g}}\right\|^2_{sq}\left\|\frac{b-\hat{b}}{\underline{b}}\right\|_{sp}\|g-\hat{g}\|^{j-2}_{(j-2)s/(s-2)}. \tag{10.12}$$

The choices $s = rj/(j+r-2)$, $p = (j+r-2)/j$ and $q = (j+r-2)/(r-2)$ give $sp = r$ and $sq = (j-2)s/(s-2) = jr/(r-2)$, and then this term is bounded above by the first term in the bias of Theorem 8.1. The strings ${}_d\varepsilon\cdots\delta_\wedge$ and ${}_\wedge\varepsilon\cdots\delta_d$ can be handled similarly; only one of the two extremes yields a factor $g-\hat{g}$, but we need to consider these strings only of length $m-1$.

This concludes the derivation of the bias. The variance is bounded by the sum of the variances of the third order estimator, and the variances of the variables $\mathbb{U}_n\chi^{(j,i)}_{\hat{p}}$ over $i = 1,\ldots,j-2$ and $j = 4,\ldots,m$, for the influence functions given in (8.6). Relative to the third order influence function, the influence functions $\chi^{(j,i)}$ add a product of projection kernels of dimension n. In the spirit of Lemma 12.4 these can be integrated out, both from the left and the right, and bounded out by the supremum on their diagonal, until only the hyperbolic, third order part of the influence function remains. The suprema on the diagonal are bounded above by a multiple of n by assumption, and the hyperbolic part can be bounded exactly as in the proof for $m = 3$. Because the higher order degeneracy of the kernels also create extra factors $1/n$ in the variances of their U-statistics (see (13.1)), the variance is no bigger than it was for $m = 3$.

11. Appendix 1: Influence functions. The main aim of this section is to prove the validity of algorithm [1]–[3] as given in Section 3 for computing influence functions. We start with a lemma that motivates the defining property of influence functions in Section 3.

LEMMA 11.1. *Let $f\colon[-1,1] \to \mathbb{R}$ and $g\colon[-1,1]\times[-1,1] \to \mathbb{R}$ be k times continuously differentiable with $g(t,t) = 0$ and $(d^j/dt^j)f(t) = (\partial^j/\partial s^j)_{|s=t}g(s,t)$ for every t and $j = 1,\ldots,k$. Then, for $j = 1,\ldots,k$ and every $u \in (-1,1)$,*

$$\frac{\partial^j}{\partial t^j}g(u,t)_{|t=u} = -\frac{\partial^j}{\partial s^j}g(s,u)_{|s=u}.$$

PROOF. The conditions show that the functions $s \mapsto f(s)$ and $s \mapsto g(s,t)$ have the same first k derivatives at $s = t$. Because also $g(t,t) = 0$, it follows

that $f(s) - f(t) - g(s,t) = o\big(|s-t|^k\big)$ as $s \to t$. By writing the remainder term in the form

$$\frac{1}{k!}(s-t)^k\Big[\big[f^{(k)}\big(t+\xi_s(s-t)\big) - f^{(k)}(t) - g_1^{(k)}\big(t+\xi_s'(s-t),t\big) + g_1^{(k)}(t,t)\Big],$$

for $f^{(k)}$ the kth derivative of f and $g_1^{(k)}$ the kth partial derivative of g relative to its first argument, we see that $f(s) - f(t) - g(s,t) = o\big(|s-t|^k\big)$ as $|s-t| \to 0$, uniformly in (s,t), by the assumed (uniform) continuity of the kth derivatives. Now the difference $f(s) - f(t)$ can also be expanded as $-\big[f'(s)(t-s) + \cdots + f^{(k)}(s)(s-t)^k/k!\big] + o\big(|s-t|^k\big)$ as $t \to s$. A similar expansion of $t \mapsto g(s,t) = g(s,t) - g(s,s)$ follows. The lemma follows by uniqueness of a Taylor expansion. ∎

Let $\chi\colon(-1,1) \to \mathbb{R}$ be m times continuously differentiable and let $t \mapsto p_t$ be a smooth map from $(-1,1)$ to $\mathcal{P}$. Assume that $\bar\chi_t^{(j)}\colon \mathcal{X}^j \to \mathbb{R}$ are symmetric functions such that, for $t \in (-1,1)$ and $j = 1,\ldots,m-1$, and for every $(x_1,\ldots,x_j) \in \mathcal{X}^j$,

$$\frac{d}{dt}\chi(t) = \int \bar\chi_t^{(1)}(x)\frac{\partial}{\partial t}p_t(x)\,d\mu(x), \tag{11.1}$$

$$\frac{\partial}{\partial t}\bar\chi_t^{(j)}(x_1,\ldots,x_j) = \int \bar\chi_t^{(j+1)}(x_1,\ldots,x_j,x)\frac{\partial}{\partial t}p_t(x)\,d\mu(x). \tag{11.2}$$

LEMMA 11.2. *Under (11.1)-(11.2) and regularity assumptions, the functions* $\chi_0^{(j)} = D_p\bar\chi_0^{(j)}$ *satisfy, for every* $j = 1,\ldots,m$,

$$\frac{d^j}{dt^j}_{|t=0}\chi(t) = \frac{d^j}{dt^j}_{|t=0}\Big(P_t\chi_0^{(1)} + \frac{1}{2}P_t^2\chi_0^{(2)} + \cdots + \frac{1}{m!}P_t^m\chi_0^{(m)}\Big).$$

PROOF. By Leibniz's rule, for every j and i,

$$\frac{d^j}{dt^j}P_t^i\chi_0^{(i)} = \sum_{j_1,\cdots,j_i}\binom{j}{j_1\cdots j_i}\int\cdots\int \chi_0^{(i)}(p_t^{(j_1)}\times\cdots\times p_t^{(j_i)})\,d\mu^i,$$

where $p_t^{(j)}$ is the jth partial derivative of $t \mapsto p_t$. Upon evaluation at $t = 0$ all terms in the sum with one of the indices $j_1,\ldots,j_i$ equal to zero vanish, by degeneracy of the function $t \mapsto \chi_0^{(i)}$. This will happen for every $(j_1,\ldots,j_i)$ if $i > j$. It follows that the right side of the lemma can be written as

$$\sum_{i=1}^{j}\int\cdots\int \chi_0^{(i)}b_0^{(i,j)}\,d\mu^i,$$

for the functions $b_t^{(i,j)}$ defined by

$$b_t^{(i,j)} = \frac{1}{i!} \sum_{j_1,\cdots,j_i>0} \binom{j}{j_1 \cdots j_i} p_t^{(j_1)} \times \cdots \times p_t^{(j_i)}.$$

We shall show that the left side of the lemma can be written in the same form.

In fact, we prove by induction on j that, for every t,

$$\frac{d^j}{dt^j}\chi(t) = \sum_{i=1}^{j} \int \cdots \int \bar{\chi}_t^{(i)} b_t^{(i,j)}\, d\mu^i. \tag{11.3}$$

Because the functions $b_t^{(i,j)}$ are degenerate, the right is unchanged if $\bar{\chi}_t^{(i)}$ is replaced by its degenerate part (in $L_2(p_t)$), and hence the lemma follows.

Let a dot denote differentiation relative to t. For $j = 1$ the identity is exactly assumption (11.1), because $b_t^{(1,1)} = \dot{p}_t$. If assertion (11.3) is true for j, then it follows by differentiation that,

$$\begin{aligned}\frac{d^{j+1}}{dt^{j+1}}\bar{\chi}(t) &= \sum_{i=1}^{j} \int \cdots \int \Big[\dot{\bar{\chi}}_t^{(i)} b_t^{(i,j)}\, d\mu^i + \bar{\chi}_t^{(i)} \dot{b}_t^{(i,j)}\Big]\, d\mu^i \\ &= \sum_{i=1}^{j} \Big[\int \cdots \int \int \bar{\chi}_t^{(i+1)} (b_t^{(i,j)} \times \dot{p}_t)\, d\mu^{i+1} + \int \cdots \int \bar{\chi}_t^{(i)} \dot{b}_t^{(i,j)}\, d\mu^i\Big],\end{aligned}$$

by (11.2). Here the function $b_t^{(i,j)} \times \dot{p}_t$ can be replaced by its symmetrization, by the assumed symmetry of $\bar{\chi}_t^{(i+1)}$. It follows that the assertion is true for $j+1$ if the $b_t^{(i,j)}$ satisfy the recursion formulas, with S denoting symmetrization,

$$\begin{aligned} b_t^{i,j+1} &= S(b_t^{(i-1,j)} \times \dot{p}_t) + \dot{b}_t^{(i,j)}, \qquad 1 < i < j+1, \\ b_t^{j+1,j+1} &= S(b_t^{(j,j)} \times \dot{p}_t), \\ b_t^{1,j+1} &= \dot{b}_t^{(1,j)}. \end{aligned}$$

The second and third recursions are consistent with the first if we set $b_t^{(j+1,j)} = b_t^{(0,j)} = 0$.

From the definition of $b_t^{(i,j)}$ we see that (note that $\dot{p}_t = p_t^{(1)}$)

$$\begin{aligned} &S(b_t^{(i-1,j)} \times \dot{p}_t) + \dot{b}_t^{(i,j)} \\ &\quad = \frac{1}{(i-1)!} \sum_{j_1,\cdots,j_{i-1}>0} \binom{j}{j_1 \cdots j_{i-1}} S\big(p_t^{(j_1)} \times \cdots \times p_t^{(j_{i-1})} \times p_t^{(1)}\big) \\ &\qquad + \frac{1}{i!} \sum_{j_1,\cdots,j_i>0} \binom{j}{j_1 \cdots j_i} \sum_{l=1}^{i} p_t^{(j_1)} \times \cdots \times p_t^{(j_l+1)} \times \cdots \times p_t^{(j_i)}. \end{aligned}$$

This can be seen to be equal to $b_t^{(i,j+1)}$. Indeed, the sum defining the latter function corresponds to the assignments of $j+1$ objects to i nonempty boxes. The two sums in the preceding display correspond to the assignments in which the $(j+1)$th object is alone in a box (i possible boxes, the other j objects distributed over $i-1$ boxes in groups of sizes $j_1, \ldots, j_{i-1}$) or is in a box with at least one other object (i possible boxes, the other objects distributed over the i boxes in groups of sizes $j_1, \ldots, j_i$). Note that the symmetrization $S\big(p_t^{(j_1)} \times \cdots \times p_t^{(j_{i-1})} \times p_t^{(1)}\big)$ can be written as an average over the i expressions obtained by placing the term $p_t^{(1)}$ before the lth factor of the product $\prod_{l=1}^{i-1} p_t^{(j_l)}$, or after the $(i-1)$th factor. ∎

12. Appendix 2: Projections. In this section we collect essential properties of projections, including representation by kernels, means and variances, and influence functions. Throughout let μ be a σ-finite measure onto some arbitrary measurable space.

12.1. *Generalities.* We call a *weighted projection* in $L_2(\mu)$ onto a closed subspace L with weight function w the map $\Pi\colon L_2(\mu) \to L$ given by

$$\Pi g = \operatorname*{argmin}_{l \in L} \int (g-l)^2 \, w \, d\mu.$$

We assume that the weight function w is bounded away from 0 and ∞, so that this map is well defined. The weighted projection is determined by: $\Pi g \in L$ and the orthogonality relationship

$$\int (g - \Pi g) l \, w \, d\mu = 0, \qquad \forall l \in L.$$

We say that the weighted projection has a *kernel representation with kernel* Π if, for all $g \in L_2(\mu)$,

$$\Pi g(x_1) = \int \Pi(x_1, x_2) g(x_2) w(x_2) \, d\mu(x_2).$$

A weighted projection is of course just an orthogonal projection onto L in the space $L_2(\nu)$ for the measure ν defined by $d\nu = w \, d\mu$, and as a kernel operator on $L_2(\nu)$ it has precisely kernel Π. On the other hand, as a kernel operator on $L_2(\mu)$ the weighted projection has kernel $(x_1, x_2) \mapsto \Pi(x_1, x_2) w(x_2)$, which includes the weight function. This ambiguity is unavoidable in our context, as we need to work with multiple weight functions, both estimated and "true" ones.

The kernel of an orthogonal projection is symmetric in its arguments. Thus with the preceding definition the "kernel of a weighted projection" is also symmetric.

Not all projections have kernels, but projections on finite-dimensional spaces do.

LEMMA 12.1. *If $e_1, \ldots, e_k$ are arbitrary linearly independent elements that span the linear subspace L of $L_2(\mu)$, then the weighted projection onto L relative to the weight function w has kernel*

$$\Pi(x_1, x_2) = \sum_i \sum_j (C^{-1})_{ij} e_i(x_1) e_j(x_2),$$

for C the $(k \times k)$-matrix with (i,j)th element $C_{ij} = \int e_i e_j w \, d\mu$.

PROOF. Because we can change measure from μ to ν given by $d\nu = w \, d\mu$, if suffices to prove the lemma for the case that $w = 1$. If $\Pi g = \sum_i \gamma_i e_i$, then the orthogonality relationships $g - \Pi g \perp e_j$ give that $\sum_i \gamma_i C_{ij} = \int g e_j \, d\nu$ for $j = 1, \ldots, k$. We can invert this system of linear equations to see that $\gamma_i = \sum_j (C^{-1})_{ij} \int g e_j \, d\nu$ for every i. Insert this into $\Pi g = \sum_i \gamma_i e_i$ and exchange the order of summation and integration to obtain the result. ∎

We view projections mainly as operators on $L_2(\mu)$, but for a number of arguments we need control of approximation errors in $L_s(\mu)$ for $s > 2$. An L_2-projection Π does not necessarily give a best approximation in $L_s(\mu)$ for $s \neq 2$, but it often gives an approximation that is optimal up to a constant. This is the case if its norm as an operator $\Pi \colon L_s(\mu) \to L_s(\mu)$ is finite. (Finiteness assumes implicitly that Π maps $L_s(\mu)$ in itself; the norm $\|\Pi\|_s$ is then by definition the minimal number C such that $\|\Pi g\|_s \leq C \|g\|_s$ for every $g \in L_2(\mu)$.)

LEMMA 12.2. *Let Π be an orthonormal projection in $L_2(\mu)$ onto a subspace L that is also contained in $L_s(\mu)$. If $\Pi \colon L_s(\mu) \to L_s(\mu)$ has bounded norm $\|\Pi\|_s$, then*

$$\|g - \Pi g\|_s \leq \big(1 + \|\Pi\|_s\big) \|g - L\|_s.$$

PROOF. The triangle inequality gives $\|g - \Pi g\|_s \leq \|g - l\|_s + \|l - \Pi g\|_s$. Since $l = \Pi l$, for every $l \in L$, the second term is bounded by $\|\Pi\|_s \|l - g\|_s$. We finish by taking the infimum over $l \in L$. ∎

One example are projections on a wavelet basis. The L_s-norm of a function is equivalent to the ℓ_s-norm of the coefficients relative to such a basis

(suitably normalized). Because the L_2-projection is the wavelet expansion truncated at a certain level of resolution, projection decreases the ℓ_s-norm of the coefficients and hence the L_s-norm of the function "up to a constant".

12.2. *Norms, means and variances.* An orthogonal projection in $L_2(\mu)$ has operator norm 1, but the square $L_2(\mu \times \mu)$-norm $\int\int \Pi^2 \, d(\mu \times \mu)$ of its kernel is equal to the dimension of its projection space.

LEMMA 12.3. *The kernel of an orthogonal projection on a k-dimensional subspace of $L_2(\mu)$ has square $L_2(\mu \times \mu)$-norm $\int\int \Pi^2 \, d(\mu \times \mu) = k$.*

PROOF. By writing the kernel in the form given by Lemma 12.1 relative to an orthonormal basis of the projection space (so that $C = I$), we find that

$$\int\int \Pi^2 \, d(\mu \times \mu) = \sum_i \sum_j \int\int e_i(x_1) e_i(x_2) e_j(x_1) e_j(x_2) \, d\mu(x_1) \, d\mu(x_2).$$

The off-diagonal elements vanish by orthogonality, while the diagonal elements are equal to 1. ∎

Typically the square norm of a projection kernel can be written as $\int \Pi(x,x) \, d\mu(x)$. In fact, the projection property $\Pi^2 = \Pi$ of a kernel operator on $L_2(\mu)$ can be expressed in the kernel as

$$\int \Pi(x_1, x_2) \Pi(x_2, x_3) \, d\mu(x_2) = \Pi(x_1, x_3), \qquad \text{a.e. } (x_1, x_3). \tag{12.1}$$

If this equation holds for every $x_1 = x_3$ and Π is symmetric, then we obtain by integration that $\int\int \Pi^2 \, d(\mu \times \mu) = \int \Pi(x,x) \, d\mu(x)$.

For simplicity of notation we assume that the kernel is such that (12.1) is valid for every x_1, x_3, in particular on the diagonal $\{(x_1, x_3) : x_1 = x_3\}$. (This is typically a null set, making this an assumption of using a special representative.) This is true in particular for the kernels in Lemma 12.1.

LEMMA 12.4. *If $\Pi_1, \ldots, \Pi_{m-1}$ are kernels of orthogonal projections in $L_2(\mu)$ that satisfy (12.1) identically, then, for any $j \in \{1, \ldots, m-1\}$,*

$$\int \cdots \int \prod_{i=1}^{m-1} \Pi_i^2(x_i, x_{i+1}) \, d\mu(x_1) \cdots d\mu(x_m) \le \prod_{i=1 : i \ne j}^{m-1} \sup_x \Pi_i(x,x) \int \Pi_j^2 \, d(\mu \times \mu).$$

PROOF. Equation (12.1) implies that $\int \Pi_i(x,y)^2\,d\mu(y) = \Pi_i(x,x)$, for every x. If $j < m-1$, then we apply this to the integral with respect to x_m of the multiple integral in the lemma, thereby turning this $m$fold integral into an $(m-1)$fold integral of the function $\prod_{i=1}^{m-2}\Pi_i^2(x_i,x_{i+1})\Pi_{m-1}(x_{m-1},x_{m-1})$. Next we bound the factor $\Pi_{m-1}(x_{m-1},x_{m-1})$ by its supremum over x_{m-1}, and are left with an $(m-1)$fold integral of the same type as before times this supremum. We repeat the argument, removing all kernels to the right of the jth kernel. Next we apply the same procedure working from the left side up, until the only remaining integral is $\int \Pi_j(x_j,x_{j+1})\,d\mu(x_j)d\mu(x_{j+1})$. ∎

The preceding results show that (under (12.1)) the square norms of (products of) projection kernels are controlled by their values on the diagonal. The following lemma shows that these values do not differ significantly for weighted projections with different weights.

LEMMA 12.5. *The weighted projections in $L_2(\mu)$ onto a finite-dimensional space L relative to the weight functions v and w possess kernels Π_v and Π_w that satisfy (12.1) identically and, for every x,*

$$\Pi_v(x,x) \le \left\|\frac{w}{v}\right\|_\infty \Pi_w(x,x).$$

PROOF. For a fixed basis $e_1,\ldots,e_k$ of L we can, by Lemma 12.1, represent the kernels as $\Pi_v(x,y) = \vec{e}_k(x)^T C_v^{-1}\vec{e}_k(y)$ for C_v the matrix with ijth element $\int e_i e_j v\,d\mu$, and similarly for Π_w. By choosing $e_1,\ldots,e_k$ to be orthonormal in $L_2(w)$ the matrix C_w can be reduced to the identity. The quotient $\Pi_v(x,x)/\Pi_w(x,x)$ then takes the form $z^T C_v^{-1} z/z^T z$ for some $z \in \mathbb{R}^k$, and it suffices to upper bound this quotient uniformly in $z \in \mathbb{R}^k$. The supremum of this quotient over z is the maximal eigenvalue of C_v^{-1}, which is the inverse of the minimal eigenvalue of C_v. Because

$$z^T C_v z = \int \Big(\sum_{i=1}^k z_i e_k\Big)^2 v\,d\mu \ge \inf_x \frac{v}{w}(x) \int \Big(\sum_{i=1}^k z_i e_k\Big)^2 w\,d\nu = \inf_x \frac{v}{w}(x)\, z^T z,$$

this minimum eigenvalue is bigger than the minimum value of v/w. ∎

LEMMA 12.6. *If $\Pi_1,\ldots,\Pi_{m-1}$ are kernels of integral operators on $L_s(\mu)$ with norms $\sup_{r/(r-1)\le s\le r}\|\Pi_i\|_s \le C$, then for arbitrary measurable func-*

tions $w_1, \ldots, w_m$ and any $r \geq 2$ (with $r/(r-2) = \infty$ if $r = 2$),

$$\Big| \int \cdots \int \prod_{i=1}^{m-1} \Pi_i(x_i, x_{i+1}) \prod_{i=1}^{m} w_i(x_i)\, d\mu(x_1) \cdots d\mu(x_m) \Big| \\ \leq C^{m-3} \|\Pi_1 w_1\|_r \, \|\Pi_{m-1} w_m\|_r \prod_{i=2}^{m-1} \|w_i\|_{(m-2)r/(r-2)}.$$

PROOF. Let M_i denote multiplication by w_i, i.e. $M_i g = w_i g$. By Hölder's inequality the left side is smaller than, for any conjugate pairs (p_i, q_i),

$$\|\Pi_1 w_1\|_{p_1} \big\| M_2 \Pi_2 M_3 \Pi_3 M_4 \cdots \Pi_{m-2} M_{m-1} \Pi_{m-1} w_m \big\|_{q_1} \\ \leq \|\Pi_1 w_1\|_{p_1} \, \|w_2\|_{q_1 p_2} \|\Pi_2\|_{q_1 q_2} \|w_3\|_{q_1 q_2 p_3} \times \cdots \\ \times \|w_{m-1}\|_{q_1 \cdots q_{m-2} p_{m-1}} \|\Pi_{m-1} w_m\|_{q_1 \cdots q_{m-1}}.$$

We finish by choosing the conjugate pairs so that $p_1 = q_1 q_2 \cdots q_{m-1} = r$ and $q_1 p_2 = q_1 q_2 p_3 = \cdots = q_1 \cdots q_{m-2} p_{m-1}$. Then the common value in the last string is $(m-2)r/(r-2)$ and the indices of the operator norms satisfy $r/(r-1) = q_1 \leq q_1 q_2 \leq \cdots \leq q_1 \cdots q_{m-1} = r$. ∎

12.3. *Approximations of weighted projections.*

LEMMA 12.7. *Let Π_w and Π be the weighted projections onto a fixed subspace L of $L_2(\mu)$ relative to the weight functions w and 1, respectively, and let M_w be multiplication by the function w. Then, for any conjugate pairs $r^{-1} + s^{-1} = 1$ and $p^{-1} + q^{-1} = 1$, any $t \leq r$, any integer $m \geq 2$, and any g,*

(i) $\big\|(\Pi_w - \Pi M_w) g\big\|_r \leq \|\Pi\|_s \|\Pi_w g\|_{rq} \|w - 1\|_{rp}$.

(ii) $\big\|(\Pi_w - \Pi) g\big\|_r \leq \|\Pi\|_s \big\|(I - \Pi_w) g\big\|_{rq} \|w - 1\|_{rp}$.

(iii) $\big\|(\Pi_w - \Pi M_w)^m g\big\|_t \leq C^{m-1} \|\Pi\|_s \|\Pi_w g\|_{rq} \|w - 1\|_{rp}^m$, *where* $p = (m-1)t/(r-t)$ *(with $p = \infty$ if $r = t$) and the constant C is the supremum of the norms of the operator $\Pi \colon L_u(\mu) \to L_u(\mu)$ over $u \in [t, r)$.*

(iv) $\|\Pi M_w g\|_r \leq \|\Pi\|_s \|\Pi_w g\|_r \big(2 + \|w\|_\infty\big)$.

PROOF. (i). The orthogonality relationships for the projections Π and Π_w imply that $\int \Pi(wg) l \, d\mu = \int wgl \, d\mu = \int w(\Pi_w g) l \, d\mu$, for every $l \in L$ and g. Because $\Pi_w g - \Pi(wg)$ is contained in L, it follows that, for every $k \in L_s(\mu) \cap L_2(\mu)$,

$$\int \big(\Pi_w g - \Pi(wg)\big) k \, d\mu = \int \big(\Pi_w g - \Pi(wg)\big) \Pi k \, d\mu, \\ = \int \Pi_w g (1 - w) \Pi k \, d\mu \leq \|\Pi_w g\|_{qr} \|1 - w\|_{pr} \|\Pi\|_s \|k\|_s,$$

by Hölder's inequality. By approximating a general element $k \in L_s(\mu)$ by a sequence in $L_s(\mu) \cap L_2(\mu)$ (truncate k by a constant and restrict it to sets of finite μ-measure that increase to the whole space), it is seen that the far left side is bounded by the far right side of the display for any $k \in L_s(\mu)$. Assertion (i) follows, because the norm $\|(\Pi_w - \Pi M_w)g)\|_r$ is the supremum of the left side over all $k \in L_s(\mu)$ with $\|k\|_s \leq 1$.

(ii). Because the function $(\Pi_w - \Pi)g$ is contained in L for any fixed g, the orthogonality relationships for Π and Π_w imply, for any function k as under (i),

$$\begin{aligned}\int (\Pi_w - \Pi)g\, k\, d\mu &= \int (\Pi_w - \Pi)g\, \Pi k\, d\mu = \int (\Pi_w - I)g\, \Pi k\, d\mu \\ &= \int (\Pi_w - I)g\, \Pi k\, (1-w)\, d\mu \\ &\leq \|\Pi_w g - g\|_{qr} \|1-w\|_{pr} \|\Pi\|_s \|k\|_s,\end{aligned}$$

by Hölder's inequality. We take the supremum over k to finish the proof.

(iii). The operator $\Pi_w - \Pi$ vanishes on L, so that $\Pi_w - \Pi M_w = \Pi - \Pi M_w = \Pi M_{1-w}$ on this space. Therefore, for $m \geq 2$ and any $t' \geq t$,

$$\begin{aligned}\big\|(\Pi_w - \Pi M_w)^m g\big\|_t &= \big\|\Pi M_{1-w}(\Pi_w - \Pi M_w)^{m-1} g\big\|_t \\ &\leq \|\Pi\|_t \|M_{1-w}\|_{t' \to t} \big\|(\Pi_w - \Pi M_w)^{m-1} g\big\|_{t'}.\end{aligned}$$

Here $\|A\|_{r \to s}$ denotes the norm of an operator $A\colon L_r(\mu) \to L_s(\mu)$, and $\|A\|_r = \|A\|_{r \to r}$. Using Hölder's inequality, we see that the norm $\|M_{1-w}\|_{t' \to t}$ is bounded above by $\|1-w\|_{tt'/(t'-t)}$, for $t' > t$.

We repeat this argument $m-1$ times with the pairs (t, t') equal to (t_i, t_{i+1}) for a sequence $t = t_1 < t_2 < \cdots < t_m = r$ such that $1/t_{i-1} - 1/t_i = (t_i - t_{i-1})/(t_{i-1} t_i) = 1/(rp)$ for every i. (We divide $[1/r, 1/t]$ in $m-1$ equal intervals of length $1/(rp)$.) This results in

$$\big\|(\Pi_w - \Pi M_w)^m g\big\|_t \leq \prod_{i=1}^{m-1} \big(\|\Pi\|_{t_i} \|w-1\|_{rp}\big) \big\|(\Pi_w - \Pi M_w) g\big\|_r.$$

Finally we apply (i) to the last term.

(iv). This is a consequence of (i) with $p = \infty$ and the triangle inequality.

■

12.4. *Influence functions.* Let L be a fixed linear space of functions contained and closed in $\mathcal{L}_2(p)$ for every p in a collection $\mathcal{P}$ of densities relative to a fixed measure ν.

LEMMA 12.8. *Let Π_p be the kernel of a weighted projection operator in $L_2(p)$ onto a finite-dimensional subspace. If the subspace and weight function w are independent of p, then for almost every (x_1, x_2) the map $x_3 \mapsto -\Pi_p(x_1, x_3)w(x_3)\Pi_p(x_3, x_2)$ is an influence function of the functional $p \mapsto \Pi_p(x_1, x_2)$.*

PROOF. The projection property gives that $\Pi_p l = l$ for every $l \in L$, which can be written as, for almost every x_2, with $d\nu = w d\mu$,

$$\int \Pi_p(x_2, x_3) l(x_3) p(x_3)\, d\nu(x_3) = l(x_2).$$

Substitute a smooth path $t \mapsto p_t$ and differentiate at $t = 0$ to conclude, with $\gamma = \dot{p}_0/p_0$ the score function of the path and $p = p_0$,

$$\begin{aligned}\int \frac{d}{dt}_{|t=0} \Pi_{p_t}(x_2, x_3) l(x_3) p(x_3)\, d\nu(x_3) \\ = -\int \Pi_p(x_2, x_3) l(x_3)\gamma(x_3) p(x_3)\, d\nu(x_3).\end{aligned}$$

Because Π_p projects onto the same space L for all p, the function $x_1 \mapsto \Pi_p(x_1, x_2)$ is contained in L for every x_2, if we use the kernel given in Lemma 12.1, and hence also the function $x_1 \mapsto (d/dt)_{|t=0}\Pi_{p_t}(x_1, x_2)$. Therefore,

$$\begin{aligned}\frac{d}{dt}_{|t=0} \Pi_{p_t}(x_1, x_2) &= \Pi_p\Big(\frac{d}{dt}_{|t=0} \Pi_{p_t}(\cdot, x_2)\Big)(x_1) \\ &= \int \frac{d}{dt}_{|t=0} \Pi_{p_t}(x_3, x_2)\Pi_p(x_1, x_3) p(x_3)\, d\nu(x_3).\end{aligned}$$

Applying the second last display with l equal to the function defined by $l(x_3) = \Pi_p(x_1, x_3)$ yields that the right side is equal to $-\mathrm{E}_{X_3}\Pi_p(x_2, X_3)\Pi_p(x_1, X_3)\gamma(X_3)w(X_3)$. ∎

12.5. *Wavelets.* An orthonormal wavelet basis of $L_2(\mathbb{R}^d)$ is given in terms of functions $\psi^v_{i,j}$ indexed by a "resolution" (or scale) parameter $i \in \mathbb{N}$, a "location" parameter $j \in \mathbb{Z}^d$, and a "dimension index" $v \in \{0, 1\}^d$ (e.g. [12, 9, 8]). Each function $\psi^v_{i,j}$ is a scaled and translated version of a fixed base function $\psi^v_{0,0}$ through

$$\psi^v_{i,j}(z) = 2^{id/2}\psi^v_{0,0}(2^i z - j), \qquad i \in \mathbb{N}, j \in \mathbb{Z}^d, v \in \{0, 1\}^d.$$

The multiresolution property of wavelets entails that for each resolution level I we can expand a function $g \in L_2(\mathbb{R}^d)$ as

$$g = \sum_{j\in\mathbb{Z}^d} \sum_{v\in\{0,1\}^d} \langle g, \psi_{I,j}^v\rangle \psi_{I,j}^v + \sum_{i>I} \sum_{j\in\mathbb{Z}^d} \sum_{v\in\{0,1\}^d-\{0\}} \langle g, \psi_{i,j}^v\rangle \psi_{i,j}^v.$$

Thus the functions $\psi_{i,j}^v$ with $i \geq I$, $j \in \mathbb{Z}^d$ and $v \in \{0,1\}^d$, with $v \neq 0$ if $i > I$, span $L_2(\mathbb{R}^d)$. We consider for each I the projection obtained by leaving out the contributions of the base functions at resolution levels $i > I$, and retaining only the projections on the functions $\psi_{I,j}^v$ with $j \in \mathbb{Z}^d$ and $v \in \{0,1\}^d$.

LEMMA 12.9. *If the generating base functions $\psi_{0,0}^v$ are bounded and compactly supported, then the kernel of the orthogonal projection in $L_2(\mathbb{R}^d)$ onto the linear span of all base functions $\{\psi_{I,j}^v : i \in J_I, v \in \{0,1\}^d\}$ whose support intersects $[0,1]^d$ satisfies $\Pi(x,x) \leq C2^{Id}$ for some constant C.*

PROOF. The kernel can be written in the form

$$\Pi(z_1, z_2) = \sum_{j\in J_I} \sum_{v\in\{0,1\}^d} \psi_{I,j}^v(z_1)\psi_{I,j}^v(z_2).$$

Here J_I includes all $j \in \mathbb{Z}^d$ such that the support of $\psi_{I,j}$ intersects $[0,1]^d$. For a fixed vector (z,z) the function $\psi_{I,j}^v(z)\psi_{I,j}^v(z)$ is nonzero only if $2^I z - j$ is contained in the support of the function $\psi_{0,0}^v$. The number of vectors $j \in J_I$ such that this is the case is bounded by a constant that depends only on the support of $\psi_{0,0}^v$. For each j the product $\psi_{I,j}^v(z)\psi_{I,j}^v(z)$ is bounded by 2^{Id} times $\|\psi_{0,0}^v\|_\infty^2$. The lemma follows. ∎

It follows by Lemma 12.5 that the kernel of the projection onto the wavelet bases viewed as subset of $L_2(\nu)$ is similarly bounded, for any measure ν with a Lebesgue density that is bounded away from zero and infinity.

In Section 8 the projection has been decomposed as a sum of projections on subspaces. Within the context of wavelet bases it is natural to choose the blocks in this decomposition equal to unions of resolution levels, so that all base functions at a given refinement level are included in the same block. To this end we choose the grids $n = k_0 < k_2 < \cdots < k_R = k$ and $n = l_0 < l_2 < \cdots < l_R = k$ defined in (8.4)-(8.5), which determine the block size, equal to dyadic numbers

$$k_r = 2^{dp_r} \sim n2^{r/\alpha}, \qquad l_r = 2^{dq_r} \sim n2^{r/\beta}.$$

This can be achieved within a factor of 2^d. The basis $e_1, \ldots, e_k$ in Section 8 can be taken equal to the functions $\psi_{i,j}^v$ for $i = 0, \ldots, I$, $j \in J_i$, and $v \in \{0,1\}^d$, with $v \neq 0$ if $i > 0$. Because there are 2^d as many functions $\psi_{i,j}^v$ at resolution level $i = i_0 + 1$ than there are at level $i = i_0$, the preceding display can be satisfied.

13. Appendix 3: *U*-statistics. For degenerate, symmetric functions $f\colon \mathcal{X}^m \to \mathbb{R}$ and $g\colon \mathcal{X}^{m'} \to \mathbb{R}$ we have

$$
\begin{aligned}
P^n \mathbb{U}_n f &= 0, \\
P^n (\mathbb{U}_n f)(\mathbb{U}_n g) &= \frac{1}{\binom{n}{m}} P^m fg, \qquad \text{if } m = m', \\
P^n (\mathbb{U}_n f)(\mathbb{U}_n g) &= 0, \qquad \text{if } m \neq m'.
\end{aligned} \tag{13.1}
$$

(If the functions f and g are not symmetric, then the second equation needs correction.) The variance of $\mathbb{U}_n f$ for a general measurable symmetric function $f\colon \mathcal{X}^m \to \mathbb{R}$ can be obtained from this formula by decomposing f in its Hoeffding decomposition

$$
f(X_1, \ldots, X_m) = \sum_{A \subset \{1, \ldots, m\}} f_{|A|}(X_A),
$$

where $f_{|A|}(X_A)$ is the orthogonal projection of $f(X_1, \ldots, X_m)$ onto the set of square integrable random variables that are measurable functions of $X_A := (X_i\colon i \in A)$ that are orthogonal to the random variables that are measurable functions of X_B for any $B \neq A$ (see e.g. [40], Section 11.4). Because the terms of this decomposition are orthogonal and each term with $A \neq \emptyset$ is degenerate and symmetric, we have

$$
\operatorname{var}_p \mathbb{U}_n f = \sum_{l=1}^{m} \binom{m}{l} \frac{1}{\binom{n}{l}} P^l f_l^2.
$$

The conditional expectation $\mathrm{E}_p\big(f(X_1, \ldots, X_m) \,|\, X_A\big)$ is the projection of $f(X_1, \ldots, X_n)$ onto the functions of X_A. The variable $f_{|A|}(X_A)$ can also be obtained by subtracting from this conditional expectation its projection onto the linear span of the functions X_B for $B \subsetneq A$. Consequently, the variance $P^l f_l^2$ is bounded by the variance of $\mathrm{E}_p\big(f(X_1, \ldots, X_m) \,|\, X_A\big)$. This yields the upper bound, still for symmetric functions f,

$$
\operatorname{var}_p \mathbb{U}_n f \leq \sum_{l=1}^{m} \frac{\binom{m}{l}}{\binom{n}{l}} \int \Big[\int f(x_1, \ldots, x_m)\, dP^{m-l}(x_{l+1}, \cdots, x_m) \Big]^2 dP^l(x_1, \cdots, x_l).
$$

For asymmetric f, this remains true if the integral in the right is replaced by the maximum value of the same expression when evaluated with the arguments of f permuted. Also note that $\binom{m}{l}/\binom{n}{l} \le (m/n)^l$, for any $l \le m \le n$.

If two measures P and Q are close, then the variances of a U-statistic under P and Q are obviously close too. The following quantifies this in terms of the order of the U-statistic.

LEMMA 13.1. *For any measurable function $f\colon \mathcal{X}^m \to \mathbb{R}$ and any probability measures $P \ll Q$,*

$$P^n(\mathbb{U}_n f)^2 \le 2m\Big(1 + \Big\|\frac{dP}{dQ}\Big\|_\infty\Big)^{2m} Q^n(\mathbb{U}_n f)^2.$$

PROOF. For $\Lambda_n = \prod_{i=1}^n (p/q)(X_i)$ the likelihood ratio of the product measures P^n and Q^n, we can write the left side as

$$\begin{aligned} Q^n(\mathbb{U}_n f)^2\Lambda_n &= \frac{1}{\binom{n}{m}^2} \sum_{|A|=m} \sum_{|B|=m} Q^n f(X_A) f(X_B)\Lambda_n \\ &= \frac{1}{\binom{n}{m}^2} \sum_{|A|=m} \sum_{|B|=m} Q^n f(X_A) f(X_B)\Pi\Lambda_n = Q^n(\mathbb{U}_n f)^2 \Pi\Lambda_n, \end{aligned} \tag{13.2}$$

where Π is the orthogonal projection in $L_2(Q^n)$ onto the closed linear span of the set of all functions of $(X_1, \ldots, X_n)$ that can be written as a measurable function of $2m$ variables only. As the X_i are independent (cf. [40], Section 11.4.), $\Pi\Lambda_n = \sum_{l=0}^{2m} \Pi_l \Lambda_n$, for

$$\Pi_l\Lambda_n = \sum_{|A|\le l} (-1)^{l-|A|} \mathrm{E}_Q\Big(\prod_{i=1}^n \frac{p}{q}(X_i) \,|\, X_A\Big) = \sum_{|A|\le l} (-1)^{l-|A|} \prod_{i\in A} \frac{p}{q}(X_i),$$

since $\mathrm{E}_Q(p/q)(X_i) = 1$. It follows that

$$\big\|\Pi\Lambda_n\big\|_\infty \le \sum_{l=0}^{2m} \sum_{k=0}^{l} \binom{l}{k} \Big\|\frac{p}{q}\Big\|_\infty^k = \sum_{l=0}^{2m} \Big(1 + \Big\|\frac{p}{q}\Big\|_\infty\Big)^l.$$

The lemma follows by substituting this bound in (13.2). ∎

James Robins
Lingling Li
Rajarshi Mukherjee
Eric Tchetgen Tchetgen
Departments of Biostatistics and Epidemiology
School of Public Health
Harvard University
E-mail: robins@hsph.harvard.edu
lili@hsph.harvard.edu
rmukherj@hsph.harvard.edu
etchetge@hsph.harvard.edu

Aad van der Vaart
Mathematical Institute
Leiden University
P.O. Box 9512
2300 RA Leiden
The Netherlands
E-mail: avdvaart@math.leidenuniv.nl